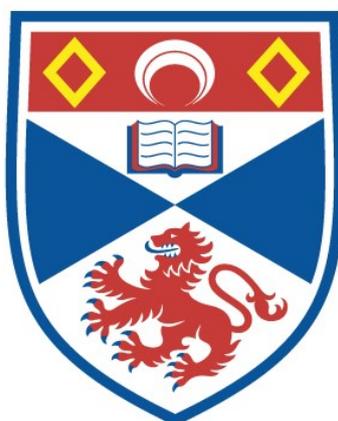


A PLURALISTIC SOLUTION TO THE RELATIONISM VERSUS
REPRESENTATIONALISM DEBATE

Paul Andrew Black

A Thesis Submitted for the Degree of PhD
at the
University of St Andrews



2018

Full metadata for this item is available in
St Andrews Research Repository
at:
<http://research-repository.st-andrews.ac.uk/>

Please use this identifier to cite or link to this item:
<http://hdl.handle.net/10023/15877>

This item is protected by original copyright

A pluralistic solution to the relationalism versus
representationalism debate.

Paul Andrew Black



University of
St Andrews

This thesis is submitted in partial fulfilment for the degree of

Doctor of Philosophy (PhD)

at the University of St Andrews

March 2018

Candidate's declaration

I, Paul Andrew Black, do hereby certify that this thesis, submitted for the degree of PhD, which is approximately 78,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for any degree.

I was admitted as a research student at the University of St Andrews in October 2013.

I received funding from an organisation or institution and have acknowledged the funder(s) in the full text of my thesis.

Date 23rd July 2018

Signature of candidate

Supervisor's declaration

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of PhD in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date 19th July 2018

Signature of supervisor

Permission for publication

In submitting this thesis to the University of St Andrews we understand that we are giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. We also understand, unless exempt by an award of an embargo as requested below, that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that this thesis will be electronically accessible for personal or research use and that the library has the right to migrate this thesis into new electronic forms as required to ensure continued access to the thesis.

I, Paul Andrew Black, confirm that my thesis does not contain any third-party material that requires copyright clearance.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

Printed copy

No embargo on print copy.

Electronic copy

No embargo on electronic copy.

Date 23rd July 2018

Signature of candidate

Date 19th July 2018

Signature of supervisor

Underpinning Research Data or Digital Outputs

Candidate's declaration

I, Paul Andrew Black, hereby certify that no requirements to deposit original research data or digital outputs apply to this thesis and that, where appropriate, secondary data used have been referenced in the full text of my thesis.

Date 23rd July 2018

Signature of candidate

Abstract

A great deal of philosophy of perception literature has been concerned with determining the fundamental philosophical account of perception. The overwhelming majority of contemporary work in this area has advocated for either a relational view of perception (broadly known as relationalism) or a representational view of perception (broadly known as representationalism). Each of these views is argued by its proponents to constitute the fundamental philosophical account of perception. These arguments are often framed in a manner suggesting that relationalism and representationalism are incompatible with one another on the grounds that if one theory explains all that we would like a philosophical theory of perception to explain, then the other theory is at best screened off as explanatorily redundant. This is known as a screening off argument, and has been utilised by both sides of the relationalism versus representationalism debate. The aim of this thesis is to demonstrate that both the utilisation of this argument in the philosophy of perception, as well as the explanatory methodology underpinning this utilisation, are misguided. This is accomplished by proposing instead that a methodology called explanatory pluralism, which holds that the best explanation of a given phenomenon is determined by what it is about that phenomenon one wishes to understand, should be applied to the debate in question. Once this plausible methodology is applied, I argue, instances that appeared to settle the relationalism versus representationalism debate decisively in favour of one view or the other instead become instances that shape the contours of a view according to which relationalism and representationalism are in fact compatible. I identify and argue for such instances, using them to support the conclusion that relationalism and representationalism are complementary explanations of perception and are, therefore, compatible.

Acknowledgements

I am grateful first and foremost to Simon Prosser for supervising, keeping me to task, offering countless helpful suggestions and observations, and for reading draft after draft of each chapter of this thesis. Thanks to Mike Wheeler for casting a fresh eye over earlier incarnations and offering invaluable feedback on them. I am grateful to audience members in attendance at the Friday Graduate Seminar in St Andrews and the 2nd Beijing Analytic Philosophy Conference at Renmin University for useful feedback and questions on several papers whose core issues feature in this thesis. I am hugely grateful to the wonderful group of people who were willing to proof-read the chapters that follow. Many thanks to all philosophers who have spoken to me about this thesis for all of the helpful discussion, suggestions and criticism throughout. Thank you also to Fiona Macpherson for introducing me to the topic of this thesis, for bringing the philosophy of perception to St Andrews (in more ways than one!), and for being a friend. Thanks also to Dr McHardy for introducing me to philosophy, and to all teachers and support staff at the University of Glasgow, Stirling University, the University of St Andrews and at UCLA who have assisted me on my journey through this strange corner of academia! Immense gratitude and love to the McDiarmids, the Fusaros, and the rest of the Luvians team – past and present – for being a St Andrews family to an academic orphan and for being saviours of my sanity more than once! Thanks also to the wonderful St Andrews Brewing Company staff – past and present – for the same reasons, and particularly the sorely missed Tom Holt for bringing everyone together. The largest thanks of all to my parents, whose unwavering support and willingness to put up with me has made this possible, and to Carley for always believing in me, and for her truly impressive amount of effort and dedication in helping me bring this all together.

Funding

This work was supported by the St Andrews & Stirling Graduate Programme (SASP).

A pluralistic solution to the relationalism versus
representationalism debate.

Paul Andrew Black

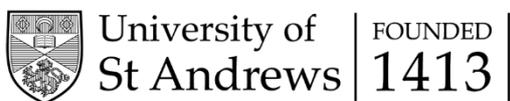


Table of contents

0: Introduction.....	9
0.1: <i>Setting the debate in context</i>	11
0.2: <i>Representationalism and relationalism: introducing the debate</i>	15
0.3: <i>'Screening off' objections</i>	18
1: Representationalism.....	22
1.1: <i>Historical context</i>	23
1.2: <i>Experiential content views</i>	26
1.3: <i>The nature of content</i>	31
1.3.1: <i>(Non-)conceptual content</i>	32
1.3.2: <i>Wide/narrow content</i>	37
1.3.3: <i>Types of content</i>	47
1.3.4: <i>Types of representation</i>	56
1.4: <i>Misrepresentation</i>	65
2: Relationalism.....	70
2.1: <i>A dialectical context for relationalism</i>	70
2.1.1: <i>Separatism</i>	72
2.1.2: <i>Qualia & Mary</i>	73

2.1.3: <i>Mary and Naïve Realism</i>	76
2.2: <i>Relationalism</i>	79
2.2.1: <i>The relational view</i>	79
2.2.1.1: <i>The direct acquaintance relation</i>	80
2.2.1.2: <i>The role of phenomenal character</i>	88
2.2.2: <i>The Object View variant</i>	90
2.3: <i>Motivations in veridical cases</i>	94
2.3.1: <i>The grounding role of experience</i>	94
2.3.2: <i>The problem of determinacy</i>	96
2.4: <i>Disjunctivism</i>	100
2.4.1: <i>Metaphysical Disjunctivism</i>	102
2.4.2: <i>Epistemological Disjunctivism</i>	120
3: <i>Explanatory pluralism</i>	129
3.1: <i>Introducing explanatory pluralism</i>	131
3.2: <i>Contextualising explanatory pluralism</i>	139
3.3: <i>Motivating explanatory pluralism</i>	143
3.3.1: <i>Explanatory pluralism in history and the social sciences</i>	144
3.3.2: <i>Explanatory pluralism in psychiatry</i>	147
3.3.3: <i>Explanatory pluralism and medical science</i>	154
3.4: <i>Explanatory pluralism in the philosophy of perception</i>	157

4: The distinct explanatory interests of relationalism and representationalism.....	166
4.1: <i>The interests of relationalism</i>	166
4.1.1: <i>Conclusive reasons</i>	170
4.1.2: <i>Demonstrative reference</i>	181
4.2: <i>The interests of representationalism</i>	186
4.3: <i>Further examples: non-veridical experiences</i>	192
5: On complementary explanations.....	202
5.1: <i>Strongly and weakly complementary explanations</i>	203
5.2: <i>The Argument from Complementary Explanations</i>	212
5.3: <i>Relationalism & representationalism as complementary explanations</i> ..	219
6: Conclusion.....	230
6.1: <i>Overcoming the Screening Off Argument</i>	231
6.2: <i>The pluralist solution</i>	235
6.3: <i>Concerns</i>	239
6.4: <i>Features</i>	243
6.5: <i>Future work</i>	244
Bibliography.....	247

0: *Introduction*

The aim of this thesis is to address and resolve a debate at the heart of the philosophy of perception. By perception, I mean the conscious experience of our environment that is informed by sensory modalities such as seeing, hearing, smelling and so on. The philosophy of perception seeks to understand this aspect of conscious experience and provide overarching philosophical theories about it. It is with regard to these theories that the debate I will be concerned with occurs.

Two theories have enjoyed discussion and success throughout the majority of contemporary philosophy of perception literature: according to the first, perception is a matter of one standing in a particular relation to the objects in the world that one perceives; according to the second, it is a matter of representing one's environment as being a particular way (what all of this means and amounts to will become clearer in due course). Call the former *relationalism* and the latter *representationalism*. Between them, the two have divided the majority of philosophers working in this area, as each claims to provide the most fundamental philosophical account of perceptual experience.

I claim that this either/or split is misguided, in that neither offers the most fundamental philosophical account of perception *simpliciter*. By applying an alternative methodology to the relationalism versus representationalism debate, I will argue that the merits offered by each view can be accommodated without precluding those of the other. This will be achieved by conceiving of each theory as offering explanations of perceptual experience and conceiving of explanations as offering answers to particular questions dictated by what it is one wishes to understand. This conception allows for a *pluralist* stance to be taken towards explanation, which (very roughly) holds that different sorts of explanation may be more effective than

others at explaining different things. By plugging certain explanations offered by relationalism and representationalism into this framework, it is possible to construct a pluralistic account that resolves this debate and – I will argue – accommodates much of the insightful work in the literature, as well as diagnosing the erroneous nature of the initial debate.

I devote the first two chapters of the thesis to providing a thoroughgoing analysis of representationalism and relationalism. Chapter 1 will focus on representationalism: what it holds, arguments in its favour, and internal debates among its proponents as to the nature of representational content. Chapter 2 provides a similar exegesis of relationalism's workings and arguments in its favour, with a particular focus on various ways relationalists have attempted to account for so-called problem cases posed by illusory and hallucinatory experiences. Chapter 3 provides a commentary on explanatory pluralism – the methodology I shall be applying to the relationalism versus representationalism debate – assessing reasons for its motivation in extraneous disciplines with the purpose of providing an argument for its motivation in the philosophy of perception. Chapter 4 evaluates side-by-side the explanations offered by relationalism and representationalism of various aspects of perceptual experience, with a view to identifying certain questions to which we seek answers that can only be answered by one theory or the other, but not both. With these cases in hand, Chapter 5 elucidates the nature and means of constructing complimentary explanations, arguing that relationalism and representationalism offer explanations that can fall within this framework, and concluding that this picture accommodates the merits of both views whilst negating the need to endorse at most one of the theories. Finally, Chapter 6 discusses the ramifications of this line of argument for the philosophy of perception, considering and addressing potential objections, and closes with some remarks on avenues for future work in this area. The remainder of this introductory chapter is devoted to providing historical context for theories of perception throughout the literature, and introducing the relationalism

versus representationalism debate with which the remainder of the thesis is concerned.

0.1: *Setting the debate in context*

It is no accident that much has been made in philosophy of efforts to understand the mental aspects to our existence. Of the array of mysteries and 'big questions' that the discipline addresses, it is the philosophy of mind that centrally treats *us*, or at least our mental workings, as the mystery to be solved – as what is to be understood. Of course, within this broad enterprise of understanding how the mind works, there are various separable aspects of the mind and its workings into which our investigative enterprises might be categorised. Of these, one that has enjoyed a great deal of attention over the years has been the philosophy of perception. Even cursory consideration of the role of perceptual experiences renders this readily understandable: it is true that not all aspects of our mental lives, as it were, are concerned with providing us with information about the external, mind-independent world and its workings. It is also true to say that a great deal of our conscious experience is dedicated to providing us with this information – it is broadly through our perceptual faculties that we enjoy success in this area.

Broadly, we might speak of the mind-*independent* world as the collection of (predominantly) physical objects and states of affairs that obtain between them irrespective of whether any subject is around to experience or act upon them. In other words, things which are mind-independent are such that were a subject to simply cease to exist – without any further consequences – this would not impact the object's existence. That is, the object would not also cease to exist. This might be contrasted with the property of being mind-*dependent*, whereby things that are correctly predicated with this term are such that, were a given subject to cease to exist, they would too – in other words: they depend on the mind of that subject for their very existence.

On the face of it, that the mind-dependent is a domain over which the mind enjoys control seems obvious: metaphysically speaking, it is at least controversial to suggest that anything mind-dependent exists on the same plane as anything mind-independent. For example, the love one feels for another, whatever else may be said about it, is not as *physically* robust as the wedding ring worn as an expression of that love. The former would (in centrally important respects) perish with the demise of both parties as a matter of necessity, whilst the latter would not. This is not to say, however, that the mind-independent and mind-dependent are wholly isolated from one another. Indeed, as the external world is thought of as being mind-independent, any information that we have about it (which might be manipulated by us mentally to our own various ends) is such that it requires some connection – one that bridges the gap between the mind-independent, external, physical plane and the mind-dependent, internal mental workings of our minds.

Perception seems a natural candidate for such a connection between the internal and external: a connection which allows us to capture information about the mind-independent world around us; information which we can then (plan to) act upon, manipulate, and so on. For one thing, then, it seems as though perception allows us to greatly enhance our epistemic states. Cohering with approaches to popular philosophical themes, perception *qua* source of knowledge is certainly a process that we have been eager to understand. Indeed, after knowledge attained via reason and intellect alone, it appears other knowledge we possess is in some way delivered to us via our perceptual faculties, be this by means of vision, olfaction, or any of the other sense modalities.

The philosophical interest in perception is not exclusively an epistemological one. Whilst there are epistemologically-minded people who concern themselves principally with what it is to know something about one's environment as a result of perceptual experience – or with precisely *how* this works – it is also true that people are occupied with the *phenomenology* of our perceptual experiences, i.e. with 'what it is like' to have them. These are

merely two out of several focal areas that are presently being investigated and theorised about in an active fashion, but which have been so for a great number of years.

Faced with this physical process that apparently allows us to gather information about the physical states of affairs and happenings in the mind-independent world, and have this information available to us mentally when planning our actions, it is only natural that we might wish to understand exactly how the process works. Throughout the years this was attempted primarily by attending to the question of exactly *what* it was that we were most directly perceiving. In other words, in seeking to understand perceptual experience, philosophers of perception turned first and foremost to address the question of what constitutes the *object of perception* in our perceptual experiences. We can see motivations for this on two fronts. In the first instance, we have it that perception appears to be the key source of much of our knowledge pertaining to the external world, yet we know also from experience that perceptual experience can be misleading. How such a fallible source of information might also be simultaneously the central foundation of such knowledge is a question of critical importance to the discipline; these are the broad concerns that motivate the so-called ‘problem of perception’, on which I shall say more below.¹

On the other hand, contemplation of a more metaphysical sort leads us quite naturally to several considerations. Firstly, veridical perception – where reality and our perception of it match – seems poised to give us knowledge *only* on the basis that our perceptions and the mind-independent objects that they are perceptions of are linked in an appropriate manner. Yet it seems clear that they cannot be one and the same thing. Take, for instance, my perception of the mountain I am looking at – this cannot be identical to the mountain itself: the mountain itself will persist if I suddenly drop out of existence, whereas my perception of it would perish with me. Additionally, one of the properties of the mountain itself is the vast amount of space it

¹ Matthen (2015: 5)

occupies, yet my perception seems to take place *within me* – something that cannot possibly be achieved by the mountain itself, as it occupies much more space than I do!

It seems, then, that the question regarding the direct object of veridical perceptions is more complicated than it may first appear. Several alternatives have been suggested as candidates and, though not the primary focus of this thesis, they merit a brief mention for reasons of historical context. Answers which are now for the most part rejected have appealed to a number of things: one example being sense-data – non-physical objects residing outside the subject, our awareness of which fundamentally accounts for perception – which are relied upon in sense-datum theory. This view is the main articulation of indirect realism, which holds that we are aware of mind-independent objects only *indirectly*, via being directly aware of an intermediary, such as sense-data.² Anti-realist views, such as Berkeley's *idealism*, according to which we are not directly perceiving mind-independent objects at all, and our perception extends only to mind-dependent objects, have also been advanced.³ Theories such as the more mentally focused adverbial view, according to which perception is a state of mind adverbially modified in some way (e.g. my seeing the tomato sauce in front of me is to be semantically described as my 'visually sensing redly') have also been proposed.⁴ For various reasons, these views have been predominantly cast aside as untenable by those engaged in the philosophy of perception literature.⁵ They do, however, serve to exhibit the number and variety of responses to the demand for understanding with respect to our perceptual experiences.

² See Russell (2001) for the traditional view, and Robinson (1994) for a contemporary version.

³ See Berkeley (2012a, 2012b, 2012c).

⁴ Initially Chisholm (1957), though Kriegel (2007) offers a contemporary adverbialist spin on representation of non-existent entities.

⁵ For reasons of space, I am unable to undergo a full-dress discussion here as to *why* these theories fell out of favour, however Brewer (2011) and Crane & French (2015) provide good discussion of this.

0.2: *Representationalism and relationalism: introducing the debate*

If the aforementioned theories are all-but-universally discarded, then what has replaced them? The answer to this is largely composed of two distinct theories of perception: *representationalism* and *relationalism*. Though a much more detailed account of each view will be provided in the chapters to follow, for now it suffices to provide the following tenets:

REP: *Perception most fundamentally consists in a subject representing the mind-independent world as being a certain way.*

REL: *Perception most fundamentally consists in a relation between the subject of the experience, and the mind-independent object(s) they are perceiving, which obtains within a particular set of circumstances.*

Despite this admittedly basic articulation, it is still correct to say that the philosophy of perception has been a two-horse race between representationalism and relationalism for the last several decades. Though this is not to say anything about how the philosophy of perception *should* have been for any length of time at all – as we shall see, I take the two-horse race mentality to be completely mistaken. I shall provide a summary of why I think this before the chapter concludes, and providing more comprehensive treatment of the issue is the overall task of this thesis, and so shall be something I attempt in the chapters that follow. For now, it will be prudent to offer some account as to why this state of affairs has obtained, irrespective of whether it ought to have or not.

There are multiple ways to accomplish this. For one, we can consider the chronology of the theories in question, in particular what they were put forward to replace. Representationalism was put forward as a response to the other theories discussed above: sense-datum theory, adverbialism and idealism, for instance. The thought was that one could get around the dubious metaphysics they employ, by positing that each experience has associated with it some content (as experience is intentional, therefore about something, and what it is about is this content), and this content is grasped by the subject by means of representation. A subject thus representing their

environment to be a certain way is – according to representationalism – what perception *most fundamentally* consists in.⁶ Conversely, relationalism was advanced to provide contemporary accommodation for the early modern intuition that the direct objects of perception just are the objects that they appear to be: i.e. we perceive objects in an unmediated manner.

For all its purported benefits, it was still thought by some that representationalism fell wide of the mark in its efforts to give the most fundamental characterisation of veridical perception (by which I mean normal, successful perception and not illusion or hallucination). In particular, it has been suggested that veridical perception most fundamentally involves a direct connection with the mind-independent objects we perceive, rather than representing them as being a particular way. Relationalism attempts to accommodate this suggestion, holding that there is a relation R of direct acquaintance that obtains between a subject S , a mind-independent object o , from a standpoint c (where c is taken to encompass the lighting conditions, angle of viewing, and other such things), and however R is cashed out is what perception *most fundamentally* consists in.⁷ Note that relationalism does not preclude the occurrence of experience involving representation, any more than representationalism precludes the subject being directly acquainted with mind-independent objects – what each denies is that the other is correct about how best to fundamentally characterise perceptual experience.

The immediate problem here, as the emphasis in the previous paragraph highlights, is that we have two distinct accounts of perceptual experience, each of which claim that their position constitutes the most fundamental account of perception. By straightforward consideration of what the words mean, we can appreciate that there can be at most one theory of perception that is *the most* fundamental. From here, the fact that most of the literature

⁶ Logue (2014: 224); not all theories of perception that fall under the banner ‘representational’ adhere to all of these commitments, but for the purposes of this introductory chapter, this is a sufficient characterisation. More elaboration on the intricacies and nuances of representationalism is provided in the following chapter.

⁷ Again, this is very concise. A much more complete exposition of the various ways relationalism has been cashed out is provided in Chapter 2.

surrounding these two sorts of view have treated the two as incompatible, has itself been taken as evidence that the two theories are incompatible. According to almost all discussion on the matter, the choice between representationalism and relationalism, and – due to the falling out of favour of other sorts of theory of perception – the choice of one’s theory of perception at all, was seemingly a one-or-the-other situation. It is this precise point that I submit is misguided. In making this claim, I join a growing list of contemporary philosophers of perception who arrive at the same conclusion, though it is fair to say that the path taken to get there varies substantially depending on whose work one reads.

In her (2011) for instance, Susanna Schellenberg defends what can be referred to as a *reconciliatory* theory of perception. According to her view, it is simply false, even on the assumption that the only plausibly correct theories to be had about perception are relationalism and representationalism, that it must be one or the other that is the correct view, as opposed to both. This is then fleshed out by way of contending that representational content and the perceptual relation of direct acquaintance are “mutually dependent” in providing a fundamental philosophical account of perception.⁸ On her view, this is because having an experience with representational content involves the employment of concepts, the possession of which in turn requires that the subject bear the sort of relation to mind-independent objects that the relationalist contends is fundamental to perceptual experience – a relation of direct acquaintance that can pick out the mind-independent objects or property instances that the concept is of.⁹

The view I shall be advancing, in contrast, agrees with the core notion that it is just false that we cannot incorporate both relational and representational views into our overarching theory of perception. In fleshing out what this amounts to, my view comes closer to Heather Logue’s, as set out in her (2014), and stresses that which of representationalism or relationalism we

⁸ Schellenberg (2011: 732)

⁹ Ibid. (732 – 733)

should be appealing to in order to best afford us an understanding of perception *depends on what it is we are trying to explain*.

Motivating the immediately preceding statement, and showing how this allows us to on some level accept both relationalism and representationalism are the central tasks of this thesis, and as such require a great deal more exposition and argument before they may be completed. I return to that task in subsequent chapters. In what remains of this one, I shall set out the type of argument which it is thought has led several to think that there is no middle-ground to be had between relationalism and representationalism, before sketching a rough overview of my response to it, and offering some concluding remarks on what the ramifications of this are, how the proposed view sits with its alternatives, and outlining in more detail what is to follow in subsequent chapters.

0.3: 'Screening off' objections

So why, then, has it been taken to be the case that representationalism and relationalism are incompatible with one another? As far as I can determine, those who assented to a position of incompatibility have been guided by reasoning that emulates that of the argument immediately below.

The Screening Off Argument:

- 1) Suppose (for the sake of *reductio*) that perceptual experiences are representational (in the sense that representationalism says they are).
- 2) Relationalism is not explanatorily redundant with respect to what we want our philosophical theory of perception to explain (a commitment of relationalism).
- 3) The fact that perceptual experiences are representational is sufficient to explain what we want our philosophical theory of perception to explain.
- 4) If perceptual experiences are representational, then relationalism is screened off as explanatorily redundant with respect to what we want our philosophical theory of perception to explain (from 3).

- 5) Relationalism is screened off as explanatorily redundant with respect to what we want our philosophical theory of perception to explain (from 1 and 4).
- 6) Contradiction (2 and 5); perceptual experiences are not representational.¹⁰

The above is an instance of a ‘screening off’ objection. There are two things to immediately note about it. The first is that there is an analogous argument that comes down on the side of perceptual experiences not being relational:

- 1) Suppose (for the sake of *reductio*) that perceptual experiences are relational (in the sense that relationalism says they are).
- 2) Representationalism is not explanatorily redundant with respect to what we want our philosophical theory of perception to explain (a commitment of representationalism).
- 3) The fact that perceptual experiences are relational is sufficient to explain what we want our philosophical theory of perception to explain.
- 4) If perceptual experiences are relational, then representationalism is screened off as explanatorily redundant with respect to what we want our philosophical theory of perception to explain (from 3).
- 5) Representationalism is screened off as explanatorily redundant with respect to what we want our philosophical theory of perception to explain (from 1 and 4).
- 6) Contradiction (2 and 5); perceptual experiences are not relational.

Both arguments presented above are *screening off objections* which, more generally, hold if there are two theories attempting to explain the same thing, and one of them explains all that there is to be explained about this, then the other becomes ‘screened off’ as explanatorily redundant *even if* it can provide an explanation to that thing as well. In

¹⁰ This is an adaptation of what Logue (2014) refers to as a ‘close cousin’ to the second part of Martin’s argument against positive accounts of hallucination (more on this in Chapter 2). I have adjusted Logue’s reconstruction such that ‘what we want our philosophical theory of veridical perception to explain’ has taken the place of the experience’s ‘epistemological, behavioural and phenomenal features’. The reason for this is that I wish to remain neutral on whether or not there is any aspect of veridical experience of which we *should* want our philosophical theory to provide an explanation that is not exhausted by the initial phrasing of epistemological, behavioural and phenomenal features.

the first argument presented above, it is concluded that perceptual experience is not representational by deriving a contradiction from the assumption that it is. Simply accepting alongside this assumption that relationalism is not explanatorily redundant, and that representationalism can explain everything we want a philosophical theory of perception to explain. The other version operates in exactly the same manner, with representationalism and relationalism swapped around, therefore reaching the conclusion that perceptual experience is not relational. Taken together, these arguments seem to suggest that one cannot accept both theories: i.e. they are incompatible.

The second thing to note is that each of the conversing formats of this argument seem to be equally valid. Consequently, since these two arguments would *directly* contradict one another, then if they are both valid, it must be the case that at least one of them is unsound. I submit that it is in fact *both* of these arguments that are unsound, and additionally contend that it is premise 3 of each argument that this lack of soundness has as its point of origin.

So I take it to be mistaken both that (i) the fact that perceptual experiences are representational is sufficient to explain everything that we want our philosophical theory of veridical perception to explain, and (ii) the fact that perceptual experiences are relational is sufficient to explain everything that we want our philosophical theory of veridical perception to explain. That is to say: I take the endorsement of the Screening Off Argument by relationalism or representationalism to be in error.

As later chapters will show, I also take there to have been a further error made on the basis of this first one. In a nutshell, this amounts to the notion that relationalists and representationalists have on multiple occasions taken the apparent inability of the opposing view to provide an explanation of some aspect of veridical perception that their view *can* explain as a decisive argument in support of their own view. In a

snappier form: proponents of screening off objections are screening off more than they ought to be. I will spell out in much more detail why I take this to be the case in due course, but in summary form the idea is that once a reconciliatory view of perception is on the table, then cases where *only one* of relationalism or representationalism can offer an explanation of an aspect of perception that we want explained should not be treated as entailing the refutation of the allegedly opposing theory.

Instead, such cases should be treated as delineating the specific commitments of that reconciliatory view that draws on each individual theory, depending on what it is about our perceptual experiences we wish to understand, and how well each theory respectively explains this phenomenon. Additionally, this type of reconciliatory view accommodates persuasive arguments made in the literature by relationalists and representationalists that their respective views are correct. Reinterpreting the dialectic in this way makes sense of why relationalism and representationalism have been viewed as incompatible, whilst resolving the debate. The persuasive arguments in favour of one view or the other, when subjected to the error of screening off as explanatorily redundant more of the opposing view than is warranted, effectively predicts the 'one or the other' mentality found in much of the literature on the topic.

1: *Representationalism*

The theory that has dominated discussion in contemporary philosophy of perception is what I will be calling *representationalism*¹¹. As shall be clear by the end of this chapter, providing a specific account of what exactly this theory amounts to is exceptionally difficult. The literature covering, in the broadest sense, content views of perception is massive and this in turn has led to the generation of a great many similar yet distinct views for which it is challenging to provide an accurate taxonomy. In order to understand what is to follow in subsequent chapters, it is worth explaining the various paths that so-called *content views* have taken within the philosophy of perception, since without doing this it is not easy to see precisely how the various views work or are to be distinguished from one another. For the exegetical purposes of this chapter, I shall use ‘content view’ as an umbrella term for views of perceptual experience that involve content in some way. Not all of these content views are articulations of representationalism, yet in discussing them, I hope the commitments of representationalism will become clear. The structure of this chapter is as follows: I will first set the advancement of content views in context; then, I shall provide an exegesis of what, in the broadest of senses, the core claims of content views are, and will draw distinctions between types of content view to arrive at a sharp definition of representationalism. Then, I provide over several sections a more thoroughgoing classificatory schema of positions that philosophers of perception have taken on issues surrounding the nature of representation and representational content. The final section discusses how representationalism deals with the so-called ‘bad cases’ of perception – illusion and hallucination – in an appealing manner.

¹¹ Block (1990) opts for the term ‘representationism’, which distinguishes the view in question from what I will call herein ‘strong intentionalism’ – the view that the phenomenal character of an experience not only supervenes on its representational content but is identical to it. I opt for the distinction I do because I imagine it easier for the reader to immediately discern what is meant, e.g., by the claim that ‘strong intentionalism is a form of representationalism’ than the claim that ‘representationalism is a form of representationism’.

1.1: *Historical context*

Historically, philosophical theories of perception have tended towards answering questions regarding what exactly it is that we most directly perceive in the course of our veridical perceptual experiences.¹² Pre-theoretically, it may seem obvious that the answer to this question is, to borrow some phrasing from Brewer, the mind-independent, physical objects that we all know and love.¹³ This permits us to say that we have knowledge of our immediate environment when we successfully perceive it, because we have *direct awareness* of it, of the sort we would lack were we hallucinating. This rudimentary view is known as *direct realism*. There are several issues with this sort of view, most prominently arguments based on non-veridical experiences – illusion and hallucination. The first of these runs as follows:

The Argument from Illusion:

P1) In an illusory experience, it seems to one that something has a quality, F, which the ordinary object supposedly being perceived does not actually have.

P2) When it seems to one that something has a quality, F, then there is something of which one is aware which does have this quality.

C1) Since the ordinary object in question is, by hypothesis, not-F, then it follows that in cases of illusory experience, one is not aware of the object after all.

P4) The same account of experience must apply to both veridical and illusory experiences.

C2) Therefore, in cases of veridical experience, one is not aware of the object after all.

P6) If one is perceptually aware of an ordinary object at all, it is in either a veridical or illusory experience.

¹² Crane (1992: 2)

¹³ Brewer (2011)

C3) Therefore, one is never perceptually aware of ordinary objects.¹⁴

As stated here, the argument from illusion seeks to establish a negative conclusion that runs contrary to the spirit of direct realism as described above. We see also in premise 2 that there is nevertheless *something* apparent to the subject in cases of illusion. This premise has come to be known as the *Phenomenal Principle*, and states that if something seems to a subject to possess a sensible quality, then there is *something* of which the subject is aware that actually does possess that sensible quality.¹⁵ A natural question that follows is precisely what the nature of this thing is, if not physical and mind-independent (as held by direct realism). Bertrand Russell proposed an answer in the form of non-ordinary, non-physical, mind-dependent objects called *sense-data*.¹⁶ The character of an illusory experience is thereby explained by the sensible qualities that these sense-data actually possess. A similar case can be made in light of hallucinatory experience, and runs as follows:

The Argument from Hallucination:

P1) When I hallucinate, I am not aware of any mind-independent, physical object.

P2) When I hallucinate, I am nonetheless aware of something.

C1) When I hallucinate, I must be aware of a mind-dependent, nonphysical mental object – a sense-datum.

P3) Experiences that are phenomenally indistinguishable are of exactly the same type, *qua* mental state.

P4) If two experiences are of exactly the same type, *qua* mental state, and one involves being aware of a mind-dependent, nonphysical object, then the other also does.

P5) For every non-hallucinatory experience there is a phenomenally identical hallucinatory experience.

¹⁴ Crane & French (2015)

¹⁵ Robinson (1994: 32)

¹⁶ See Russell (2001)

C2) All perceptual experience, hallucinatory and non-hallucinatory, involves awareness of a mind-dependent, nonphysical object – a sense-datum.¹⁷

Ultimately, Russellian sense-datum theory fell out of favour due to a number of objections raised against it. For one thing, the metaphysical status of a sense-datum is dubious; for another, implausible consequences of the Phenomenal Principle upon which these arguments rely are drawn out by Armstrong, who generates a phenomenal sorites paradox using the notion of colour samples. We are asked to consider three colour samples: C_1 , C_2 , and C_3 ; these are such that (with respect to colour) C_1 appears exactly similar to C_2 , C_2 appears exactly similar to C_3 , but C_3 does not appear exactly similar to C_1 . By the Phenomenal Principle, this entails the existence of corresponding sense-data S_1 , S_2 and S_3 , which are such that (again, with respect to colour) S_1 is exactly similar to S_2 , S_2 is exactly similar to S_3 , and S_3 is *not* exactly similar to S_1 . Yet such an arrangement, it is argued, violates the alleged transitivity of exact similarity with respect to a given property, in this case colour. It is plausible to say that *apparent* colours need not be transitive – limitations in the human visual system make this quite likely – but the Phenomenal Principle entails that sensible properties of sense-data surely are just the sort of properties which exact similarities between ought to be transitive.¹⁸

Despite the apparent failure of Russellian sense-datum theory, there are other avenues to pursue from this point: the direct acquaintance that Russellian views claim that we have with sense-data, for instance, may be argued to obtain between subjects and mind-independent objects *without* a sense-datum-like intermediary (as in contemporary relationalism).

Alternatively, as will be the focus of the remainder of this chapter, we might focus on perceptual representations in order to articulate our philosophical theory of perception.

¹⁷ Macpherson (2013: 12 – 13).

¹⁸ Armstrong (1993: 218 – 219)

1.2: *Experiential content views*

Issues with other theories of perception led to the advancement of what may be labelled, in the broadest of terms, as *content views*. Outlining several options for the core commitments of such views is the primary task of this section; first, though, I shall say a little on why it might be thought that perception is representational, as this will stand us in good stead for what is to follow.

In the first instance, it may seem *prima facie* obvious that experience is intentional, in the same way that beliefs, desires and the like are: experiences seem to be *about* things. Considering the case of a belief, it is clear that “there is a way the world could be which would make the belief true and a way the world could be which would make the belief false.”¹⁹ In other words: the belief has *accuracy conditions* – when the content of my belief is that there is a white lamp on my desk in front of and to the right of me, then the accuracy conditions for this belief are met if there actually is, in the world, a white lamp on my desk in front of and to the right of me. Some beliefs can be more coarse or finely grained than others, and it is possible that the accuracy conditions may become more stringent if my belief is of a finer grain (e.g. that the lamp is ceramic, or is of the sort used to treat Seasonal Affective Disorder), and these accuracy conditions will not be met if the world does not accurately reflect the content of these finer-grained beliefs, even if the original accuracy conditions on the more straightforward, coarse-grained belief about the white lamp and its position would be.

Some proponents of the content view apply something like the foregoing analysis of beliefs’ accuracy conditions to perceptual experience, generating a view according to which experiences are intentional, and thereby according to which “all visual perceptual experiences have contents.”²⁰ The idea is that the contents of one’s experience constitute accuracy conditions for that very experience: if, as the case may be, I have an experience of a whiteboard on

¹⁹ Macpherson (2011: 2)

²⁰ Siegel (2010: 334)

the wall to my right, then (at least part of) the content of my experience is *that* there is a whiteboard on the wall to my right, and this experience will be accurate (or veridical) iff there exists in the world a whiteboard on the wall to my right. Otherwise, my experience is inaccurate, by way of being to some degree either illusory or hallucinatory.²¹ As Siegel points out, however, all that is established at this stage is that experiences possess accuracy conditions – nothing has been done to argue that these accuracy conditions are fit to serve as the content of experience. What the content theorist really wants, after all, is to uphold the claim that all perceptual experiences involve *contents*, not merely accuracy conditions.

Identifying views that have upheld this claim and assessing arguments in their favour is the major task of the remainder of this section. It will first be prudent to introduce some terminology, and in so doing make clear key distinctions between what I am taking to be three core versions of content view.²² Logue distinguishes between three ‘flavours’ of view: Mild, Medium and Spicy. I shall now explain in turn the respective commitments of each of these.

The Mild Content View consists in nothing more than an acceptance of the following two claims, which Logue contends form the bare minimum that one is committed to in holding that experience has content: for any perceptual experience *E*,

- (i) there is a proposition associated with *E*, and
- (ii) this proposition captures the way things perceptually appear to the subject in virtue of having *E*.²³

Being a fairly weak claim, the Mild Content View is silent on a key point often endorsed by theorists who accept content views. Since the view holds merely that there is a proposition *associated* with *E*, one could endorse the

²¹ Macpherson (2011: 4)

²² I stress again that some, but not all, of these content views are rightly called ‘representationalism’, as I am using the name. This will become clear below.

²³ Logue (2014: 222)

view without holding that experience is a propositional attitude. It can be denied on the Mild Content View, in other words, that the subject experiencing *E* perceptually *represents* the proposition associated with *E*.²⁴ Thus, the proposition is relevant to the theorist attempting to characterise the experience, but is not such that the subject of the experience bears any distinctively experiential psychological relation to it.²⁵ To put it another way, the Mild Content View can be held even whilst denying that the experience has any *experiential content*.²⁶

With the Mild Content View introduced, the other two views can be defined with relative ease. Starting with the Medium Content View, this is simply the preceding claims (i) and (ii), together with the following claim:

- (iii) perceptual experience consists in the subject perceptually representing her environment as being a certain way.²⁷

In terms of differences between the Mild and Medium Content Views, what most obviously differs is that the former involved no commitment whatsoever to the subject perceptually representing their environment as being a particular way and is instead merely a commitment to an *association* between an experience and a propositional content.²⁸ The Medium Content View, by the addition of claim (iii), is committed to holding that some experience *E* consists in the subject perceptually representing this propositional content, and thereby *does* hold that experience has representational content.

Nevertheless, the Medium Content View remains silent on whether or not the experience *fundamentally consists* in this perceptual representation. As Logue puts it, “[t]o say that perceptual experience fundamentally consists in personal level psychological feature *x* is to say that it has some or all of its

²⁴ Ibid. (223)

²⁵ Ibid.

²⁶ For clarity, *experiential content* is an umbrella term designed to capture all content of an experience, that is: content that the experience itself possesses (whether it is veridical or otherwise).

²⁷ Ibid.

²⁸ Ibid.

other personal-level psychological features ultimately in virtue of x ,” thus one could accept the Medium Content View whilst denying that this perceptual representation is what ultimately provides our philosophical account of perception.²⁹

Finally, the Spicy Content View is essentially an acceptance of that latter point on which the Medium Content View remains silent, and so involves a commitment to claims (i) – (iii) above, as well as to the following:

- (iv) perceptual experience fundamentally consists in the subject perceptually representing her environment as being a certain way.³⁰

This fourth claim entails that perceptual experience necessarily involves representation, and holds that it is via this representational process that we can articulate our philosophical account of perception.

With these three ‘flavours’ of content view introduced, it is possible to clarify some terminology that will feature in what follows. As I shall use the term, *representationalism* is the name given to any view which endorses above claims (i) – (iii) *at the least*. By that, I mean that representationalists may (and often do) endorse (iv) as well – such views are also representationalist. Views which endorse (iv) as well, and hold that representation is *all* that is fundamental to our philosophical account of perception, I shall call *austere representationalism* (in view of the fact that they do not appeal to explanatory resources outside of representation and representational content).³¹ Probably the most popular contemporary articulation of (often austere) representationalism is *intentionalism*.

Intentionalism holds at its core the claim that representational content determines phenomenal character (here understood as the Nagelian ‘what it

²⁹ Ibid.

³⁰ Ibid. (224)

³¹ I follow Schellenberg (2014) in this naming convention. The austerity in question is to be explained by the minimal reference to any theory other than representationalism (such as relationalism). The converse will apply later to *austere relationalism*. In the context of the representationalism versus relationalism compatibility debate, distinguishing austere forms of the theories from versions of them incorporated into the overall reconciliatory view I later advance will be useful. I take the austerity terminology to facilitate this.

is like'-ness of an experience); that is, a supervenience relation holds between the two such that there can be no difference between the phenomenal character of two distinct experiences without a difference in representational content.³² Intentionalists of all stripes adhere to this claim, yet many disagree as to how precisely it is to be cashed out. Many of these differences do not matter for my purposes, but one that does draws a distinction between *strong intentionalism* and *weak intentionalism*. Weak intentionalism is precisely the view I describe above – holding that the relation between representational content and phenomenal character is one of supervenience; strong intentionalism makes the stronger claim that the relation instead is one of identity.

I take it that the Mild Content View – that each experience *e* is merely *associated* with some content – is at the very least *prima facie* plausible.³³ However, for the relationalism versus representationalism debate that is my main focus, it is not as relevant as the others, as it can be held by (at least almost) all parties to the debate. For the purposes of addressing that debate, it is worth noting as Siegel does that “[when philosophers of perception] criticise the idea that experiences have contents, their criticism is best understood as directed at [the Spicy] Content View according to which experiences are *fundamentally* structured as a propositional attitude.”³⁴ It is important to realise, as Siegel’s words suggest, that screening off objections only stand to render representational and relational theories incompatible provided they are each making a claim to be *the* fundamental characterisation of our perception. It is this sort of content view that is most relevant to my overall purposes in what follows, and which I previously identified as falling under the banner of representationalism.³⁵

³² Byrne (2001: 204)

³³ See Siegel (2010) and especially Logue (2014) for arguments to this effect.

³⁴ (2010: 363 – 364); emphasis mine.

³⁵ Representationalism is sometimes used as an alternative name for what I have called ‘strong intentionalism’, according to which an identity relation obtains between representational content and phenomenal character. This has been acknowledged as creating an unhelpful confusion (c.f. Macpherson (2009: 507)). Confounding things further, I believe my usage goes against the grain, but I believe utilising this terminology makes the

To reiterate, as I am phrasing things, representationalism is simply the view that perception has representational content, and that this plays a key role in our philosophical account of perception. This comes apart from the claim that it is all there is to the construction of such an account. The latter view – that perception is fundamentally to be explained by representational content, and nothing else – is called *austere representationalism*. One of the main goals of this thesis is to argue against it. I take up this task in Chapter 4 onwards. Now, I consider further the nature of the representational content that is at the heart of the theory.

1.3: *The nature of content*

The task for this section is to provide some idea of the different positions that philosophers of perception have taken on the further particulars of representational content. For the purposes of keeping things broadly relevant to my overall enterprise, I am not discussing in what follows the mere association thesis of the Mild Content View. The reason for this, alluded to above, is that the overall project of this thesis is to make progress in the debate as to whether or not content views are compatible with relational views of perception, and it seems overwhelmingly plausible that the Mild Content View *is* compatible with relationalism. This is true even if the proponent of the Mild Content View accepts all steps of the Screening Off Argument on the side of relationalism. Even if the content in question provides no explanatory insight whatsoever because all that needs to be explained is explained by something else, this is perfectly compatible with the experience in question being *associated* with some content or other.³⁶

overall discussion to follow far clearer, insofar as it centrally pertains to the compatibility of the representational and relational views of perception.

³⁶ A similar argumentative move is made by Siegel with respect to an experience's accuracy conditions when she contends that "the claim that experiences can be accurate or inaccurate can be true even if the assessability of an experience for accuracy is less fundamental than its status as a 'good' case of perception." (2010: 342)

The first point to be raised about the nature of experiential content is that talking about the content of experience is not the same as talking about the content of a belief generated on the basis of that experience.³⁷ To clarify: assuming the intentionality of both experiences and beliefs, it follows by definition that each is *about* something. The point being made here is that even if the respective contents of a subject's experience and the belief that they are disposed to form on the basis of that experience both end up being contents to the effect 'that *p*', for example, they are not the same thing – one is the content of an experience, the other is the content of a belief formulated on the basis of undergoing that experience. Insofar as the two are distinct, what is being centrally discussed in this section is the *experiential* content, not the content of the belief generated on the basis of that experience.

1.3.1: (Non)conceptual content

A prominent debate as to the nature of representational content is whether it is conceptual or nonconceptual. It is this distinction and debate that I shall focus on presently, though it is worth noting that my aims at this point are purely exegetical: I am introducing the debate, as well as highlighting some reasons in favour of conceptual or nonconceptual content put forward in the relevant literature. What I am *not* attempting here is the advancement of my own view on the content of experience as it figures into my views on perception (in particular viz. reconciliatory theories thereof).

Defining each of the views in the conceptualism/nonconceptualism debate seems a sensible place to start. From McDowell, we have that “conceptualism does not take concepts merely to impose a top-down constraint on the range of permissible experiential contents. They are intimately involved in the production of that content.”³⁸ There are two claims here. The first is that the

³⁷ See Macpherson (2011) for further discussion on this point.

³⁸ Wright (2015: 181)

concepts possessed by a subject impose restrictions on the representational contents their perceptual experiences can permissibly have; the second – which explains the first – is that the concepts possessed by the subject play an integral part in generating the representational content of the experience(s) in question. If, for instance, one construes perceptual experience as being representational in the same vein as belief is representational, one might think that the content of experience integrally involves “an exercise in our grasp of concepts” in just the same way as our thoughts and judgements do.³⁹ Nonconceptualism, meanwhile, can be characterised in terms of the specification of correctness conditions, the idea being that “while concepts would have to be exercised in providing a (theory-relative) canonical statement of a given content’s correctness conditions, having an experience with that content does not require possessing or deploying any of those concepts.”⁴⁰ The key point here lies in the second part of the quotation – nonconceptualism is a rejection of the view that states that in order for a subject to have an experience with representational content, the subject in question must exercise (and therefore possess) concepts relevant to that content.⁴¹ More specifically, there is a distinction between *content* nonconceptualism and *state* nonconceptualism: the two are related, though the former is a denial of the claim that an experience’s representational content is comprised by its concepts, whilst the latter denies that in order to be in a certain representational state, one must possess the concepts reflective of that state’s representational content; conceptualism, meanwhile, tends to be a view that accepts both of the claims that these distinct versions of nonconceptualism deny.⁴²

With these working definitions of conceptualism and nonconceptualism in place, I turn now to examining the merits of each view, starting with

³⁹ Campbell (2014: 45)

⁴⁰ Wright (2015: 181)

⁴¹ Where ‘relevant’ is taken in at least the sense required for providing a complete specification of the accuracy conditions applicable to the experience in question.

⁴² Ibid. (182)

conceptualism. Perhaps the most familiar argument in favour of conceptual content stems from the idea that perceptual experience is linked to our beliefs, knowledge and reasons for action. As Wright puts it:

*If perceptual experience has the same sort of conceptual content as our thoughts about the world we can literally believe what we see, see what we want and so forth. In such a circumstance, it seems unproblematic to hold that perceptual experiences are capable of both justifying beliefs and feeding their contents into our decision-making and action-planning processes.*⁴³

The point being made here is that conceptual contents make for a much smoother story as to the relationship holding between perception and reason. This is a weaker claim, however, than that made by Brewer which looks to actively refute nonconceptualism, namely: that reasons positively *require* conceptual contents.⁴⁴ More elaborately, he holds that a subject has reasons for believing something iff they are in a conceptual state, where a conceptual state is understood as being a mental state which is such that it “has a representational content that is characterisable only in terms of concepts which the subject [themselves] must possess and is of a form which enables it to serve as a premise or the conclusion of a deductive argument, or of an inference of some other kind (e.g. inductive or abductive).”⁴⁵ The thought here seems to boil down to acceptance of the claim that reasoning requires the possession and deployment of concepts, together with the view that we want a neat story as to how exactly it is that perception is linked to reasoning. If we accept Brewer’s story that reasons require conceptual contents, then it seems as though concept possession and deployment runs the whole way through the chain from perception to rational action, with reasoning and beliefs as intermediaries.⁴⁶ Despite being different processes, the conceptualist thought goes, the link between

⁴³ Ibid. (183)

⁴⁴ (2002: 149)

⁴⁵ Ibid.; though Brewer has since become a relationalist, maintaining that these considerations constitute an objection to non-conceptual content. (2017: 219f)

⁴⁶ “We can make sense of an experience being a subject’s *reason for* believing something about the state of the world, conceptualists say, only if there are conceptual relations between the experience and the belief.” Ibid. (185)

experience and belief justification can occur precisely because the same conceptual capacities are operative in both, and these impose an order on our sensory deliverances of the sort that is arguably required for thought and experience to be related rationally.⁴⁷

A basic consideration in favour of nonconceptual content comes in the form of imagining a subject observing two colour samples that are near-identical shades of red, just distinct enough that the subject's sensory system can pick this up. For instance, samples which would be directly adjacent to one another on some colour organisational system, yet can be distinguished between despite the subject's ignorance of the colour organisational system, would be qualitatively distinguishable despite the subject's lack of some salient concepts.⁴⁸ If there is a difference in phenomenology, and thereby in phenomenal character, then at least according to intentionalism there should be a difference in representational content. If there is a difference in representational content, then it is not clear how the conceptualist might explain this since, by hypothesis, the subject lacks the concepts to account for the difference in contents.

Additionally, there is the consideration that conceptualism – given its commitment to the claim that having states with representational content requires possession and deployment of concepts – seems to spit out the curious result that infants and non-human animals that are incapable of concept possession are similarly incapable of undergoing states with representational content.⁴⁹ When coupled with the assumption that perception is representational, this seems to implausibly suggest that they are incapable of perception.

As I have stated, I do not take a side on this matter, as my present aims are wholly expository, but I do think that the devil will be in the details. To that end, the putative arguments offered above in favour of conceptual and

⁴⁷ Ibid.

⁴⁸ Ibid. (181 – 182)

⁴⁹ Ibid. (188)

nonconceptual content respectively may fall foul of the mistake that is taking the available options in the debate to be either the view they advance or the view to which they take themselves to be opposed, with no middle ground to be had.⁵⁰ A point counting against this interpretation can be seen by considering again the biconditional, laid out above, that a subject can have reasons for believing something iff they are in a conceptual state. Brewer advances this as an objection to nonconceptual content, yet it is not clear that one may go as far as this. What appears to be ruled out is the *absence* of conceptual content. What is much less clear is that the presence of nonconceptual content is ruled out: as Wright observes, “[n]onconceptualism doesn’t rule out states with conceptual content; it just denies that the latter exhaust all representational content-bearing experiential states.”⁵¹ A conceptual state, recall, is one which possesses *a* representational content that is characterisable only in terms of concepts possessed by the subject and is of the form that can be utilised in reasoning of some description. There is certainly room in logical space to claim that an experience involving its subject being in a conceptual state, as Brewer defines such a state, may also have nonconceptual content, particularly if one considers that perceptual experience – and therefore perceptual processing – is something which occurs over time, so it is not even the case that such a suggestion implies a subject possessing an experience that has conceptual and nonconceptual content *simultaneously*.⁵²

Though nothing that has been said in this section decisively settles the conceptual/nonconceptual content debate, I hope to at least have accomplished (i) elucidating precisely what each of these terms means in the context of the philosophy of perception, and (ii) highlighting some standard arguments made in favour of either sort of content. I have also suggested

⁵⁰ A mistake not dissimilar to that made by many relationalists and representationalists. I discuss the perils of *this* particular mistake at length in chapters to follow, once I have fully introduced and discussed each of the purportedly incompatible theories of perception.

⁵¹ (2015: 182)

⁵² See Wright (2015: 191 – 193) for discussion of empirical support for the sort of middle ground position whereby perceptual processing involves both conceptual and nonconceptual content over time. Peacocke (1992: 90-91) also supports the view that nonconceptual content has “relations” to certain conceptual content, and suggests that the two form something of a “holism”.

that matters may not be quite so black and white, however nothing in what follows hangs on this suggestion.

1.3.2: Wide/narrow content

Another set of options regarding the nature of representational content concerns whether or not it is – to use the terminology commonly deployed in the literature – wide or narrow content. Simply put, the view that representational content is narrow is the view that it supervenes on the internal configuration of the subject, and does not rely on the mind-independent world for its constitution. On this view, if two subjects are identical with respect to all brain states that do not involve a relation to their respective environments, then those subjects have in common all representational contents of their experiences.⁵³ Conversely, those who posit that representational content is wide reject this supervenience claim, and accept that the mind-independent physical world enters the explanatory tale of content generation at some point.

Considerations regarding whether representational content is wide or narrow are numerous.⁵⁴ Perhaps the area where this debate rages most fiercely, though, is that of intentionalism; the reason for this is that intentionalism closely ties the notion of representational content to that of the phenomenal properties of experience – the Nagelian ‘what it is like’, or phenomenal character. The latter can additionally be taken to be either entirely ‘in the head’ (phenomenal internalism) or else as somehow world-involving (phenomenal externalism). Commitment to the weak intentionalist supervenience thesis suggests, on the grounds of parsimony, that one should be of the view that phenomenal character and representational content are either both wholly in the head or are both best explained by reference to matters outside of the subject, in the mind-independent world;

⁵³ Block (1996: 20)

⁵⁴ See Chalmers (2004) for extensive discussion on many of these. For the purposes of providing one example, I confine discussion here to a more contemporary development in the debate.

obviously, the strong intentionalist identity thesis leaves no room whatsoever in logical space for disagreeing with this contention. Whether strong or weak intentionalist, this linking of representational content with phenomenal character leads to fierce debate concerning whether or not content is wide or narrow, and I spend the remainder of this section introducing and responding to Adam Pautz's argument against the view that content is wide.

Pautz's argument is designed to target any form of phenomenal externalism, but it is relevant to the wide/narrow content debate if a commitment to intentionalism is plugged in. He contends that doing this results in something akin to an intentionalist position whereby the representational states enjoyed by the subject of an experience in some way track the instantiation of properties occurring in the mind-independent physical world. This tracking plays a key role in the generation of the representational content of the subject's experiences – a view he calls *tracking intentionalism*.⁵⁵

Pautz contends that tracking intentionalism is vulnerable to what he calls *coincidental variation cases*. These are hypothetical yet physically possible cases in which a pair of subjects track the same physical property, yet intuitively differ in phenomenal experience: this flies in the face of tracking intentionalism – the same property is tracked, and thereby represented, yet there is a phenomenal difference. Pautz's offers the following as a coincidental variation case for gustatory experience:

[S]uppose Yuck and Yum belong to different species that evolved in separate environments containing some berries. Now you might suppose that Yuck is an actual human – me or you – and Yum is some hypothetical creature. Or you might suppose that Yuck and Yum both belong to hypothetical, human-like species. It does not matter. In any case, the berries are extremely poisonous to Yuck. By contrast, in Yum's environment, the berries are a very important food source, since other food sources are scarce. So Yum's species evolved immunity to the berries. In addition, when Yuck and Yum taste the berries, their taste systems

⁵⁵ Pautz (2013: 240 – 241)

*undergo radically different ensemble activation states (spatiotemporal neural patterns).*⁵⁶

In the context of the example, ensemble activation states are distinct patterns of neurons that fire when a subject undergoes particular taste sensations. Consequently, two subjects undergoing different ensemble activation states should entail a difference in their respective taste experiences. So, supposing that poison dart frogs are extremely poisonous to Yuck and Yum, we can imagine that Yuck's ensemble activation state after eating the berries is akin to his ensemble activation state upon eating a poison dart frog; conversely, Yum's ensemble activation state after eating the berry is more akin to his ensemble activation state after eating a nice banana. Additionally, we can imagine that there are differences in their behavioural responses, for example Yuck vomits and withdraws violently from the berries, whilst Yum rubs his belly, and so on. However, this difference in behaviour does not alter the fact that the physical (i.e. chemical) properties instantiated by the berries are identical, and therefore that Yuck and Yum's respective experiences track these same properties.

Here is Pautz again:

*Despite these differences, we can stipulate that Yuck and Yum are similar at the receptor level. Indeed, we can stipulate that, when they taste the berries, the postreceptor ensemble activation patterns in their taste systems, although different, optimally track the very same complex chemical property of the berries, C. This chemical property C will likely be a disjunctive property, because many different combinations of chemical properties can produce the same response in the taste system. So I am stipulating that their ensemble activation states track the same disjunction of chemical properties C, the very one with which tracking intentionalists and other objectivists about taste would identify the taste perceived by Yuck and Yum.*⁵⁷

In sum, then, the case of Yuck and Yum is set up so that the tracking intentionalist is committed to the taste experiences of Yuck and Yum being

⁵⁶ Ibid. (255)

⁵⁷ Ibid.

phenomenally identical, despite differences in ensemble activation states and behaviours, since by stipulation they are tracking the same physical properties of the berries. There is also no difference in their sensory systems that would account for a phenomenal difference between their experiences.

There are two possible conclusions to draw from the case:

Same Experiences: Yuck and Yum have phenomenally identical experiences when eating the berries.

Different Experiences: Yuck and Yum have phenomenally distinct experiences when eating the berries.

Pautz argues that Different Experiences is more plausible, due to the subjects' contrasting behaviours and ensemble activation patterns, and we should therefore accept this conclusion.⁵⁸ Since tracking intentionalism is committed to Same Experiences, it is tracking intentionalism that should be rejected.

The moral of the Yuck and Yum case generalises into *the internal dependence argument*, which runs as follows:

1) If tracking intentionalism is true, then in *every* possible coincidental variation case, the right verdict is Same Experiences.

2) But it is much more reasonable to suppose that, at least in *some* coincidental variation cases, the right verdict is Different Experiences (call this *internal dependence*).

3) So, tracking intentionalism is (probably) mistaken.⁵⁹

The argument is valid, and Pautz appears to think it sound even against views advocating a complex tracking relation – one which may feature further *relata* involved in the subject tracking a physical property. This is because of the 'same tracking' stipulation: regardless of how complex the

⁵⁸ Ibid. (256)

⁵⁹ Ibid. (268); I maintain the parenthesised qualifier from the original work. I take it to be included because of the appeal to reasonable supposition in (2).

tracking relation is, it is stipulated that in coincidental variation cases, the subjects are tracking the same property.

Despite the apparent effectiveness of the internal dependence argument against tracking intentionalism, an appeal to Millikan's biosemantics permits a view both compatible with tracking intentionalism's tenets *and* capable of sustaining the Different Experiences conclusion, thus challenging Pautz's argument. To achieve this, the view must satisfy two criteria:

- (i) it is compatible with tracking intentionalism's tenets, and
- (ii) it can uphold the Different Experiences conclusion in coincidental variation cases without violating (i).

In the rest of this section, I will outline biosemantics and explain how it satisfies (i). In the next section I will articulate a more refined form of biosemantics, arguing that this satisfies (ii). Finally, I conclude with some more general remarks about my argument and its potential to be generalised for views besides mine.

Broadly, Millikan's view is that representational systems can be thought of as consisting in a producer of the representation, and a consumer of the representation - it is the latter of these that she takes to determine representational content.⁶⁰ This is because without the consumer understanding the representation, the representation cannot be deemed to have been successful. The producer must therefore produce a sign that corresponds to the state of affairs represented in a manner intelligible to the consumer. Further, the view claims that this successful correspondence and conveyance is a "normal condition for the proper functioning of the consumer device as it reacts to the representation."⁶¹

The most fundamental kind of representations, according to Millikan, are known as 'pushmi-pullyu' representations: when they occur in a normally functioning system, these are both descriptive, in that they inform the

⁶⁰ Millikan (1989: 286)

⁶¹ Ibid. (287)

subject enjoying them as to what is the case, and directive, in that they also inform the subject what they should do, given this presentation of the world.⁶² For example, consider the dance that honeybees perform. This dance maps on to the location of nectar (descriptive), as well as the direction and duration of flight required of on-looking bees to reach it (directive). Consequently, had either variable been altered sufficiently, the nectar would have to be in a different location relative to the observing bees in order for the representation to be successful.⁶³ If we take the dancing bee to be the producer of the representation, and take an on-looking bee to be the consumer of the representation, then in more abstract terms, what Millikan takes to happen here is that "the producer produces a sign that will be true or satisfied only if it maps on to some affair or affairs [. . .] in the world in accordance with certain 'semantic' rules. These are rules of correspondence between signs and world affairs that have been instantiated in the past when the consumer and producer or their ancestors have succeeded in performing their cooperative function(s)."⁶⁴ Thus, we have it on this theory that descriptive and/or directive representations are interpreted by the consumer according to certain rules, which are established by the successful conveying of information within this system by producer and consumer, or by their evolutionary ancestors. Hence, the view has it that a given representation fulfils a particular function, and is integrated into the system of representation *because* it fulfils that function.⁶⁵

In virtue of being a naturalised theory of representation, I take biosemantics to be straightforwardly compatible with tracking intentionalism's

⁶² Millikan (2009: 396)

⁶³ Ibid. (398)

⁶⁴ Ibid. (397 – 398)

⁶⁵ This is not to suggest that the content of intentional signs (the bee dance in the above example) is determined by what particular task it (in conjunction with its consumer(s)) normally accomplishes, for one sign may serve to accomplish various tasks in different circumstances; rather, the content is determined by the fact that "whatever inferences its consumer makes when functioning properly, the result will be (nonaccidental) true belief or successful action only if it actually represents according to a certain correspondence rule." (Ibid. (400)) This serves to highlight the importance of the correspondence rules to this theory, for it is of course possible that true information may be conveyed to the consumer accidentally (i.e. not via the relevant correspondence rule) in a manner which benefits the consumer, yet this will not constitute the system operating in its normal conditions.

commitment to objective reductionism. Since it takes representational content to be determined widely, acceptance of any view on which the phenomenal character is thereby also determined widely (such as intentionalism) results in objective reductionism. The point here is not that biosemantics *entails* objective reductionism, merely that it is compatible with it.

Interpretation of the intentional sign in accordance with the relevant correspondence rules is, I take it, an internal process; it is also one that involves being sensitive to the intentional sign that maps on to properties instantiated by external stimuli. Biosemantics satisfies the second tenet of tracking intentionalism, by requiring that successful representation occurs only when the consumer undergoes an internal representational state that can be described as tracking the instantiation of intentional properties in external stimuli.

Admittedly, since biosemantics is a theory of representation rather than one of the phenomenology of experience, it is not obvious that a proponent of the view is bound to accept the intentionalist thesis that constitutes the third tenet of tracking intentionalism. A biosemantics on which the intentional properties of experience are explained in the manner described above, yet the experience's phenomenal properties are explained by something else, is conceivable. This does not preclude satisfaction of condition (i) above, however – biosemantics does not entail intentionalism, but is compatible with it. The neutrality on phenomenal properties (and their connection to intentional properties) leaves intentionalism open as a view to be accepted in conjunction with the commitments of biosemantics, which is all that (i) requires. Call the biosemantic view that *does* endorse intentionalism *biosemantic intentionalism*. This view would satisfy all three tenets of tracking intentionalism, thereby fulfilling criterion (i) identified above, and I will now argue it satisfies criterion (ii).

Millikan's biosemantics views representation as a product of evolutionary selection: representation fulfils a teleological function and is selected for

because it does so. Organisms track physical properties because doing so benefits the organism. Let us return to Yuck and Yum. Yuck - according to biosemantics - has evolved from ancestors for whom the berry was poisonous, and thus has come to represent the berry as being bitter/disgusting by way of tracking chemical property C in the berry. Yum, conversely, is such that the berry is *not* poisonous to him and has, by stipulation, evolved from ancestors who generated an immunity to the berry due to it being an important food source. Comparing this with the bee example from earlier, we could say that the bee dance itself parallels the berry flavour in this instance, as both convey information to the consumer of the representation. Further, there is a descriptive element in the tracking of the berry's chemical properties that mirrors the tracking of the nectar's location, as well as a directive element that suggests to Yuck and Yum to respectively avoid or consume the berries, which tracks historical benefits this course of action has had for their ancestors. This mirrors the tracking of the observing bees' location relative to the nectar.

For Yuck and Yum (and their evolutionary ancestors), there is benefit in tracking chemical properties in the berry. By Pautz's 'same tracking' stipulation, both track chemical property C of the berry. Pautz's attack on tracking intentionalism claims that this ensures the intentional properties of both Yuck and Yum's experiences are identical, meaning that - due to intentionalist commitments - their gustatory experiences must be phenomenally identical.

It is not clear that the intentional properties of Yuck and Yum's experience of the berries *are* identical, however. One way that this conclusion could be established is by demonstrating that Yuck and Yum, for all that has been said, are not tracking the same physical properties. Consequently, there can be a difference in the intentional properties of their experiences, and Different Experiences can be maintained without violating tracking intentionalism. This is what has been called the *pluralist response*, and accounts both for subjects tracking two different physical properties in the same object, and for one or both subjects tracking a property not

instantiated by the object (thereby undergoing an illusion). This is not the response I wish to make here. Pautz stipulates that whatever is required for Yuck and Yum to be tracking the same physical properties obtains, and so it will not be effective. Instead, I advocate for biosemantic intentionalism, and claim that this view correctly delivers the Different Experiences conclusion. On this view, the Tracking tenet will be a necessary though not a sufficient condition for identity across the intentional properties of Yuck and Yum's respective experiences.

If the intentional properties of an experience can be defined as what that experience is *about*, then it seems that Yuck's experience is - on Millikan's view - in some way about the fact that the berry *is poisonous to him*. Avoiding the poison is precisely the beneficial effect for Yuck (and his ancestors) that caused tracking to occur in this case. Conversely, the benefit for Yum (and his ancestors) is that the berry is an important food source since other food is scarce, and so it makes sense to say that in some respect this is what his experience is about for him. If Yuck and Yum differ in what is beneficial to them in this in this respect, then it is arguable that what their respective experiences are about (on biosemantic intentionalism) differs, and therefore the intentional properties of their experiences differ. This accounts for the Different Experiences verdict in a manner that doesn't violate any intentionalist commitments.

Furthermore, this is not the pluralist response discussed above that Pautz's 'same tracking' stipulation comes ready to rebut, and so the 'same tracking' stipulation cannot be used in this way, especially not "without going into the details of Millikan's sophisticated consumer-based theory of representation."⁶⁶

Thus, I am assenting to the view that Yuck and Yum are tracking the same physical properties, but disagree with Pautz's contention that this alone entails that Yuck and Yum have experiences with the same intentional properties. I am contending (in line with biosemantics) that it is not merely

⁶⁶ Pautz (2013: 293)

tracking that determines the intentional properties of an experience; rather, the function that the tracking fulfils for the organism, and appropriate usage of correspondence rules, both play vital roles in doing this. On biosemantic intentionalism, then, it is possible for two subjects to track the same physical properties whilst having experiences which differ in representational content so long as the *correspondence rules* that determine this content are different. Given the evolutionary story told in the Yuck and Yum example, it is plausible that the correspondence rules in operation for the subjects' respective communities are different. This accounts for a difference in the intentional properties of the subjects' experiences, thus allowing Different Experiences, and the three tenets of tracking intentionalism, to be maintained. Hence, criterion (ii) from the preceding section is satisfied.

Accepting that biosemantic intentionalism is a form of tracking intentionalism falsifies the conditional claim that if tracking intentionalism is true, then in *every* coincidental variation case, the right verdict is Same Experiences. Since this conditional constitutes premise (1) of Pautz's internal dependence argument, the argument is refuted. I do not see any obvious rejoinder available to the proponent of the internal dependence argument that can help re-establish (1). The basis for accepting it was that subjects in all such cases had experiences with the same intentional properties, and therefore had phenomenally identical experiences. In light of the above point about differing correspondence rules, assuming these play a crucial role in determining intentional properties, any treatment of coincidental variation cases such as Yuck and Yum which concludes that the subjects have experiences with the same intentional properties forgoes any claim to being realistic - a feature that Pautz is eager to maintain in order to mitigate the non-actual nature of his coincidental variation cases.⁶⁷ The commitments of biosemantic intentionalism suggest that the role of correspondence rules in determining the intentional properties of experience cannot realistically be ignored. If the foregoing is correct, then premise (1) of the internal

⁶⁷ Ibid. (254)

dependence argument cannot be established, and therefore the argument fails.

The purpose of this section has been to highlight some contemporary moves being made in the debate regarding wide/narrow content, and the role this plays in the larger question as to the nature of representational content, particularly when intentionalism is accepted. I have discussed Pautz's internal dependence argument as it purportedly weighs against wide content intentionalism, and offered a response that I take to show that such a position is not refuted by Pautz's argument. As with vast majority of this chapter, my aim here has been purely expository, and to show that none of the variety of views I have discussed are radically implausible.

For all that has been said in this section, it is worth observing that narrow content is a live option as things stand. One objection to narrow content which I do not intend to discuss comes from Block, who advises against this move for intentionalists, as narrow content intentionalism must embrace functionalism. This, he argues, brings with it several concerns to which the intentionalist cannot respond due to their qualia scepticism.⁶⁸ In the following section, I address a debate on type of content, which looks to have numerous ramifications for proponents of narrow content. I wish to emphasise that my aims in this section and the next are purely expository. I survey arguments about the matters discussed merely to provide the reader an overview of the relevant discussion points. Nothing in the chapters to follow relies on these debates being resolved in a particular way.

1.3.3: Types of content

Another issue surrounding representationalist theories of perception concerns the type of content such views appeal to: a stance must be taken on what exactly constitutes this content. One influential response to this

⁶⁸ I address the notion of qualia further in Chapter 2 (see §2.1.2). See Block (1996: 20 – 23) for further discussion on this functionalism-based objection.

issue holds that the contents of a given representation simply consist in the worldly extensions of the of its constituents: if S represents that Freya is charismatic, then according to this view the contents of the representation are Freya herself, and the property of being charismatic, as these are the extensions of the words 'Freya' and 'charismatic'.⁶⁹ On this reading, one constituent of the representational content (Freya) is the referent of a singular term (i.e. one that refers to an individual), rather than descriptions that pick out different things depending on what in the world satisfies them, whilst the other (charismatic) is a property this individual instantiates. Such contents have often been called *Russellian* contents, "because at one time Russell advocated the view that singular terms contribute only their referents to the proposition expressed by sentences in which they occur (or utterances of such sentences)".⁷⁰

Russellian contents, in virtue of being purely extensional – i.e. reducible to worldly objects and their properties – are *wide* contents. If one is an intentionalist, then one's commitments arguably preclude one from holding that the determination of phenomenal content and that of phenomenal character can come apart.⁷¹ The wide nature of the content on this view entails that if two subjects are internal duplicates, then the content of their respective experiences must be determined by the same worldly extension of the proposition associated with the content. A tension between Russellian content and phenomenal internalism was observed with regards to colour experience by Ned Block via his spectral inversion cases.⁷² A similar tension is observed by Brad Thompson, where we are asked to consider representational content of experiences of size and shape – i.e. *spatial content*.⁷³ According to the Russellian account of content, two subjects with phenomenally identical spatial experiences must be such that their

⁶⁹ Siegel (2016)

⁷⁰ Ibid.

⁷¹ On an intentionalist position, phenomenal content is taken to be that representational content that determines phenomenal character (Prosser (2011: 476)). Prosser (2016: 93) provides an argument for the claim that, on intentionalism, phenomenal content and phenomenal character are either both narrow or both wide.

⁷² See Block (1996).

⁷³ See Thompson (2010).

experiences attribute the same properties to the objects perceived.

Thompson provides two thought experiments designed to show that this is a problematic commitment.

The first of these thought experiments is Doubled Earth. The setup we are asked to consider is a world that is identical to ours in every respect, except that everything is twice the size on Doubled Earth as it is on our planet.⁷⁴ We are asked to consider a pair of subjects: Oscar and Big Oscar, who are functional duplicates enjoying experiences – on Earth and Doubled Earth respectively – of objects which are counterparts to one another.

Consequently, if the object Oscar is experiencing is ten metres away from him, then the object Big Oscar is experiencing is twenty metres away from him.⁷⁵ Since everything else is held constant, including the scale of spatial *relations* obtaining between various objects (such that Doubled Earth is a ‘to scale’ duplicate of Earth), the thought goes that Oscar and Big Oscar are phenomenal duplicates.⁷⁶

Were this to be the case, then in the event that Oscar and Big Oscar had phenomenally identical experiences which were such that Oscar’s is caused by an object ten metres away whereas Big Oscar’s is caused by something twice that distance away from him, the argument runs that we would not be inclined to say that either of their experiences failed to be veridical (as privileging one over the other would merely be chauvinism). If each experience is veridical, then each must truly attribute a different spatial (physical) property to the objects respectively perceived. This, based on the commitments of Russellian content identified above, entails that if one

⁷⁴ Other differences between Earth and Doubled Earth are limited to those necessary, in light of the key difference in sizes, for experiences on the two planets to be functionally identical to one another. A reduction in time lag for light travelling distances, or signals being sent to or from the brain, for instance.

⁷⁵ So long as one keeps constant the unit of measurement: the object Big Oscar sees is twenty ‘earth metres’, as it were, away from him, though measuring in the Double Earth version of metres (which, of course, are twice the size of their Earthly counterparts) will yield that they are still ten metres away in that system. The important point is that the *space itself* occupied by everything on Doubled Earth doubles.

⁷⁶ Thompson (2010: 155 – 156)

wishes to uphold intentionalism, then representational content pertaining to size is not Russellian.⁷⁷

A corresponding thought experiment, based on the preceding one, is also constructed so as to motivate a similar conclusion regarding experiences of shape. The setup is extremely similar, except that Doubled Earth has been replaced with 'El Greco World' – a world in which everything is twice the *height* of its relevant counterpart on Earth. More specifically, it is stretched vertically, relative to the centre of the planet: a ball rolling on earth is such that all points on the surface (circumference) of the ball are equally far from the radius; a ball rolling in El Greco World, on the other hand, is such that it not only looks like an ellipse, but there are two points on the ball's surface that are furthest from the radius – one touching the ground and the other on the top – and these constantly change as the ball rolls. Comparably, we can return to Oscar, who is such that the distance between the top of his head and soles of his shoes is six feet, whether he is standing up or lying down. Stretched Oscar of El Greco World, on the other hand, is such that the distance between the top of his head and soles of his shoes is six feet whilst he is lying down, but *twelve* feet when he is standing up.⁷⁸ As Thompson puts it, the vertical 'stretch' in El Greco World is *plastic*, rather than rigid and one-off.⁷⁹ Again, it is submitted that Oscar and Stretched Oscar are veridical perceivers phenomenal duplicates, despite the counterpart objects they are respectively perceiving having different shapes. If this is so, then for similar reasons as above, Russellianism about the content of visual experiences with regards to shape *also* seems to be false.

Taken together, the thought experiments appear to show that upholding intentionalism comes at the cost of conceding spatial content is not Russellian if phenomenal character supervenes on the subject's internal configuration. An argument is made by Simon Prosser in his (2016) that a similar tension holds for the content of our temporal experiences. We are

⁷⁷ Ibid (157 – 158)

⁷⁸ Ibid. (176 – 177)

⁷⁹ Ibid. (176)

asked to consider *Slow Earth*: an atom-for-atom physical duplicate of Earth, except that everything that happens on Slow Earth happens at half the rate it does on Earth, perhaps because Slow Earth lies in a region of space where for some reason all matter is governed by laws of physics isomorphic to those of the rest of the universe, but is slowed down by half.⁸⁰ With regards to one's experience, it is observed that if one were to enter the slowed region of space, one would not notice everything slowing down, as *all* processes – including one's own bodily and particularly neural functions – would slow down together at the same rate; one would be disposed to make all of the same observations one would have made outside the slowed region, one would just make them more slowly.⁸¹

With the now familiar Twin Earth-like scenario in place, an argument analogous to those above is made for our experience of time. We are asked to consider Horatio and Slow Horatio, who are internal duplicates in every respect, and whose lives are identical in every way, except that Slow Horatio occupies Slow Earth, and therefore his life unfolds at half the rate of Horatio's. If the spatial scenarios plausibly involved phenomenal duplicates, then it seems equally plausibly that Horatio and Slow Horatio are phenomenal duplicates: all events that befall them are identical, though those that befall the latter occur at half speed, yet do so to scale – phenomenally, there should be no subjectively detectable differences for either subject.

If they are truly phenomenal duplicates, then intentionalism and phenomenal internalism together entail that Horatio and Slow Horatio's phenomenal content is identical. Given the setup of the case, it follows from this that “the phenomenal content of rate-of-change phenomenology cannot consist in objective rates of change,” because by hypothesis the respective *objective* rates of change on Earth and Slow Earth differ, yet the rate-of-change *phenomenology* – and the phenomenal content that determines this –

⁸⁰ Prosser (2016: 87 – 88)

⁸¹ *Ibid.* (88)

for Horatio and Slow Horatio is the same.⁸² This constitutes an analogous argument to Thompson's, but for the temporal case. Consequently, it is arguable that temporal content is not Russellian if phenomenal character supervenes on the subject's internal configuration. If one also finds the spatial case plausible, then the case can be made that this conclusion applies to the content of both spatial *and* temporal cases.

Though appeal to phenomenal externalism is one option available for those wishing to reject the preceding arguments, it is not the only one: there are alternative views on the type of content that representational content is that are capable of upholding phenomenal internalism and intentionalism together. One popular alternative to Russellian content presented in the literature is *Fregean* content, which holds that rather than being wholly reducible to objects and their properties, the contents of experience are instead composed of the *modes of presentation* of those extensions.⁸³ The key notion here is that objects (and the properties instantiated by them) may have multiple modes of presentation: they may, broadly speaking, appear to the subject in different ways or occupy different roles in their reasoning.⁸⁴ A popular philosophy of language example that illustrates this concerns Hesperus and Phosphorus – also called both the morning star and the evening star – which are names that denote the planet Venus. A subject might rationally accept the proposition 'Hesperus is the morning star' whilst rejecting or remaining neutral on the proposition 'Phosphorus is the evening star', if they are ignorant of the fact that Hesperus is Phosphorus, and hence of the fact they are simply different modes of presentation of the planet Venus. In such cases, the propositions the subject assents to (or not) cannot differ in their truth value, as they refer to the same object, but the different modes of presentation play different roles in the subject's reasoning, such as listing Hesperus in the set of celestial objects visible in the morning but not

⁸² Ibid. (95)

⁸³ Siegel (2016)

⁸⁴ Ibid.

in the set of celestial objects visible in the evening (or vice versa with Phosphorus).⁸⁵

Fregean contents of experiences play a similar role, and can attribute a difference in representational contents of two experiences of the same property (or properties) to a difference in modes of presentation: i.e. to a difference in *how* the object is represented. Views according to which experiences possess both a Russellian wide content *and* a Fregean narrow content – where the latter is the mode of presentation of the former – have been argued to avoid the unpalatable consequences of the above arguments.⁸⁶ The Fregean content of experiences of the colour green, on such views, can be something like ‘the property that typically causes experiences with phenomenal character *G in me*’ and this then determines the Russellian content: whatever the relevant extension is.⁸⁷ Importantly, this allows for Inverted Spectrum cases, where for me that extension might be something I would call green, whilst for an alternative subject the very same Fregean content might pick out something that I would call red: in such cases, the experiences are alike in Fregean narrow content, but differ in Russellian wide content.⁸⁸

Consider again Horatio and Slow Horatio. Suppose that each is staring at their microwave waiting for their food to defrost. They both experience the duration of the wait: it is ten minutes (their microwave timers confirm this fact), and is equally dull for each of them. By hypothesis, their phenomenal experiences are identical, yet the objective duration of Horatio’s wait is half that of Slow Horatio’s. The view of content described immediately previously can make sense of this, as both subjects will have experiences with identical Fregean contents of something like ‘the duration that brings about temporal experiences of ten minutes elapsing *for me*’, which account for their identical experiences’ phenomenology whilst functioning as modes of

⁸⁵ Ibid.

⁸⁶ See Chalmers (2004), Thompson (2010).

⁸⁷ Chalmers (2004)

⁸⁸ Prosser (2016: 94)

presentation for distinct Russellian contents: the different objective temporal durations that they pick out.⁸⁹

The Fregean intentionalist view of content described above is not the only alternative to straightforward Russellian content, however, and others have been sought. One reason for this is that the Fregean view of content makes ineliminable reference to phenomenal character of experience, and is therefore non-reductive – a fact which has not appealed to those with physicalist aspirations.⁹⁰ An alternative approach that does not do this is defended in Simon Prosser's (2011), which contends that phenomenal contents are *affordance relations*, where these are understood to be relations obtaining between a subject and objects in their environment, and which pertains to the possibilities for causal interaction between the two.⁹¹ Considering the content of the utterance 'x is heavy', the thought is that in describing x as heavy, I am not describing a property merely of x, but am – strictly speaking – describing a relation between x and myself that pertains to the possibility of causal interaction between us.⁹² In particular, I am articulating that x is going to be difficult for me to lift. I am featured inalienably, though only implicitly, and this point generalises: according to Prosser's view, contents specified as simply 'heavy' or 'poisonous', for instance, involve the subject of the experience as an unarticulated constituent.⁹³ My experience of the property of weight, then, correlates on this view with the 'heaviness' phenomenology for a given subject, which in turn is cashed out in terms of an affordance relation defined in terms of the causal powers of both the subject and object in question.⁹⁴

⁸⁹ This example I am using is borrowed from Prosser, whose view differs from the Fregean content view. This particular paragraph has recounted the Fregean account of the example; I discuss the particulars of Prosser's view below.

⁹⁰ Ibid.

⁹¹ Prosser (2011: 479); see also his (2016), where the terminology moves away from 'affordance relation' to 'SEF relation' (Subject-Environment Functional Relation), but on the whole the view of content defended is the same, save for the former discussing spatial content whilst the latter discusses temporal content.

⁹² Ibid. (480)

⁹³ See *ibid.* (480 – 481)

⁹⁴ Ibid.

This notion of affordance relation is taken to deal with spatial content of the sort addressed by Thompson and described above. Since ‘big’, ‘small’, ‘near’ and ‘far’ can all be analysed as relations between the subject experiencing those properties and the objects to which they are predicated, in terms of possibilities for causal interactions between the two. Returning to Oscar and Big Oscar, we can imagine a pair of scenarios to illustrate the point. Suppose that each is out walking in woodlands and a stream crosses each of their respective paths. We can suppose further that the opposite bank is objectively two Earth-metres away, such that Oscar would be correct to report it as being two metres away, whilst Big Oscar would be correct to report it as being one metre away (because of the scaled-up nature of Doubled Earth). In considering whether or not the bank is too far away to simply step over the stream, we can imagine that Oscar would conclude that it is, whereas Big Oscar would reach the opposite conclusion. Analysing ‘far’ (more specifically, ‘too far *for me*’) as an affordance relation in the manner described above therefore makes sense of this spatial case. Temporal cases are treated with a similar analysis: since objectively the different durations occurring in Horatio’s experience of ten minutes and Slow Horatio’s experience of ten minutes play identical roles in each of their lives, the objective duration of the former stands in the same relation to its subject as the objective duration of the latter does to its subject. All possibilities for causal interaction remain identical between the two, and this functional identity accounts for the phenomenal sameness of their respective experiences, despite the objective durations of the experiences differing.

I cannot attempt to provide decisive arguments for or against any of the types of content identified in this section, and stress again that my aims for this chapter are purely expository. I take the discussion to have identified several different directions that those wishing to decisively settle what type of content representational content is have taken, with some considerations for why they hold the views that they do offer, and with examples provided to assist in understanding these. The view to be advanced in later chapters remains largely ecumenical on these issues and provides further discussion

in instances where it does not. I shift focus now from types of content to types of representing.

1.3.4: Types of representation

The final distinction I wish to discuss in this part of the chapter is not so much about the nature of representational content *per se* so much as it is about the way in which it is represented. It is the distinction between allorepresentation and autorepresentation that is due to Charles Travis. I shall first briefly define each of these terms, before moving on to discuss Travis's argumentative utilisation of them, and thereafter highlighting a response to his argument due to Brogaard. I close with some remarks on where the discussion of this subsection leaves content views viz. the nature of representation and representational content.

Travis's distinction focuses around the thought that representational content is commonly thought to involve taking things in one's environment to be 'thus and so', that is to say that it is part of the content of a subject's experience that their environment - what they are perceiving of it, at least - is a particular way. With this in mind, Travis first defines *autorepresentation* as simply "representing something *to oneself* as so."⁹⁵ By way of clarification, this is contrasted with either producing or being aware of something that represents something as being a certain way, whereby one can then accept this production or object of awareness as thus representing.⁹⁶ The reason for the 'auto' in autorepresentation is that the subject is - as the above emphasis highlights - representing things to themselves as being a particular way. *Allorepresentation*, on the other hand, is precisely that which was being contrasted with just previously, and *does* centrally involve being aware of or producing something that represents things as being some way, such that the subject is then open to accepting or rejecting this.⁹⁷ The key

⁹⁵ Travis (2004: 61); emphasis mine.

⁹⁶ Ibid. (60)

⁹⁷ Ibid. (61)

difference upon which the distinction hangs lies in the answer to the question ‘what is doing the representing?’. In both cases, it is the subject of the experience who is being represented to (i.e. they are the recipient of the represented information). In the case of autorepresentation, it is the subject that is actually representing things as so; in the case of allorepresentation, it is something extraneous to the subject that is representing something *to* them – they are the passive recipient, who can either accept or reject what is being represented (that things are thus and so).

With these varieties of representation in play, I turn now to reconstruct Travis’s argument to the effect that, as Brogaard puts it, “no mental state represents, except in a derivative sense,” where this is understood to mean “that it’s not perceptual experience that represents but the experience together with certain higher-order epistemic states; viz. those that involve a commitment on the part of the agent to things being thus and so.”⁹⁸ So as to remain as consistent as possible with what is to follow, here is Brogaard’s reconstruction:

P1) If perceptual experiences represent [in either sense], then they represent the way things perceptually appear/look to be.

P2) If perceptual experiences themselves represent, then they represent independently of the agent’s particular epistemic states (i.e. their rational decisions, beliefs, etc.) [which is to say: they allorepresent].

P3) There is no unique way that things perceptually appear/look to be, independently of the agent’s particular epistemic states identified in (P2).

C1) Hence, perceptual experiences are not [allo]representational.⁹⁹

The first premise looks to be safe from criticism, at least from anyone who wishes to resist Travis’s conclusion – as we saw from the arguments of Logue and Siegel in favour of the content views, something like (P1) is an integral part of such theories. Additionally, if ‘the way things perceptually

⁹⁸ Brogaard (2017: 274); I use ‘varieties of representation’ somewhat loosely here. I will take care to specify which sort I mean in hereafter reconstructing and discussing Brogaard’s response to his argument.

⁹⁹ Ibid.

appear/look to be' implies a uniqueness of how things seem *for the subject of the experience*, then the argument is valid.¹⁰⁰ The second premise finds support in the observation that the content of experience and the content of beliefs formed on the basis of that experience are two distinct things, together with the contention that certain aspects of early perceptual processing are not subjectively accessible: both potentially support the contention that (at least conscious) perception involves immediate awareness, as opposed to mediated awareness or any sort of inference.¹⁰¹ Motivation for the third premise comes from two contrasting ways in which the word 'looks' (or, alternatively, 'appears') might be utilised. One of these – call it the *comparative* use – is the sort being utilised when I say 'Mark looks like his sister'.¹⁰² This refers in particular to the appearance of things (say, visually). In contrast, there is another (arguably less strictly correct) usage of the word 'looks' which might be called the *epistemic* use of looks: as in 'it looks as though the answer is 4' where this usage is not independent of the speaker's epistemic states in the way that the comparative use is.¹⁰³ The key point for motivating Travis's third premise is that the comparative use is compatible with many different ways in which things may look alike, and therefore fails to pick out a particular way things look, and therefore cannot fix representational content, unlike the epistemic use of 'looks' which – although saddled to the subject's epistemic states such as beliefs and so on – *does* successfully pick out a particular way things look, and therefore can fix representational content.¹⁰⁴ If this latter point is correct, then it seems as though there is no way that things can look a particular way to a subject that is independent of that subject's higher-order epistemic states – precisely the

¹⁰⁰ Ibid.

¹⁰¹ Ibid.

¹⁰² I alter the example slightly from Travis's original 'Pia looks like her sister' so as to illustrate that the comparative 'looks' operates over comparisons of the form 'x looks like y with respect to z'. Presumably Mark and his sister do not look identical – perhaps, for instance, they merely have the sort of similarity with respect to facial bone structure that one frequently sees with non-identical siblings, including those of distinct biological sex.

¹⁰³ Ibid. (274 – 275)

¹⁰⁴ Ibid. (275)

claim being made in the third premise. Together with the validity of the argument, it seems therefore as though the conclusion follows and is true.

It is the third premise to which Brogaard provides a response.¹⁰⁵ Her response ultimately boils down to an observation that humans have evolved to have the sort of visual system that can calculate, e.g., constancies in size, colour and shape according to what she (after Pylyshyn (1999)) calls “perceptual principles”.¹⁰⁶ Importantly, these perceptual principles operate independently of an individual’s beliefs or other epistemic states, and so do not imply that perceptual experience is cognitively penetrated (although they are not inherently incompatible with such a position either); rather, they depend not on rational capacities of a subject, but instead on the visual system that that subject has evolved to have – in this way, they are “constitutive of perceptual experiences in beings like us.”¹⁰⁷ As these perceptual organising principles are capable of calculating a particular way that things are to appear to the subject, and operate independently of that subject’s epistemic states, it seems that they constitute the required counterexample to Travis’s third premise, and thus his overall argument is blocked.

I do not take the above presentation to decisively settle the matter, but it is worth noting that if Brogaard’s argument goes through, then perceptual representation is not ruled out in the manner that Travis suggests. In any event, this matters less for my purposes, as I instead think that an earlier step in Travis’s argument is mistaken. He argues that if perceptual representation occurs at all, then it does so in the form of allrepresentation, as it cannot do so in the form of autorepresenting. I think

¹⁰⁵ This may strike some as surprising, given the apparent potential to block the second premise via reference to cognitive penetration (see Siegel (2012)). This looks like precisely the sort of case that would undermine the premise in question inasmuch as it involves beliefs influencing the nature of perceptual experience. Yet if one relies on this as a response, as Brogaard (2017: 279) observes, then Travis’s overall point goes through – there is no way that things perceptually seem to the subject independently of higher-order epistemic states.

¹⁰⁶ Ibid. (279)

¹⁰⁷ Ibid. (280)

this statement is misguided, and shall provide below his reasons for making the claim, before providing an argument in support of my contention.

We saw above that the second premise of Travis's argument is supported by the observation that perceptual experience and the beliefs formulated on its basis are separate. In addition, it has been widely accepted that the former operates to some degree independently of the latter, which is to say that doxastic or other epistemic states cannot influence perceptual experience without limit.¹⁰⁸ A classic example of this is the Müller-Lyer illusion, whereby two lines of equal length are placed next to one another, and one is embellished with inward-pointing hashes at either end, whilst the other is embellished with outward-pointing hashes at either end. One can apparently believe wholeheartedly that two lines are of equal length, yet nevertheless experience one (the one with outward-pointing hashes at either end) as being longer than the other. Further, claims Travis, autorepresentation does not permit one the option to accept or reject the content of the representation – to autorepresent something just is to accept it.¹⁰⁹ If perception is a source of information, then autorepresentation falls short of it, on Travis's view, precisely because it is not a source of anything but, in the manner described, is a matter of having already accepted the content of the autorepresentation as being so – it is a matter of *registering*, or at least of presuming to.¹¹⁰ Since perception is plausibly a source of information – a source that one can accept or reject – it is concluded that perceptual representation cannot involve autorepresentation.

Travis presents the two latter points as independent, though on reflection it is hard to see a meaningful distinction: that autorepresentation is a matter of (presuming to) register information seems just to be a function of taking the content of that information to be how things are. In any case, I take his conclusion to be misguided. For one thing, the initial justification that

¹⁰⁸ See Fodor (1983). Much of the contemporary literature on cognitive penetration treats the findings of this work as a benchmark.

¹⁰⁹ Travis (2004: 62)

¹¹⁰ *Ibid.*

perceptual experience operates at least somewhat independently of beliefs formed on the basis of (or about) that experience provides us with a means to accept or reject the deliverance of perceptual experience, even if the latter has already been accepted – or registered – in our experience via autorepresentation. More generally, it is simply not obvious that the capacity to accept or reject what the senses tell us is actually exercised at an earlier point in the perceptual process than the subject forming judgments on its basis. Consider again the Müller-Lyer illusion: we saw that our experience of the lines will be of them as differing in length, but rejection of this sensory deliverance is a matter of us *believing* that they are the same length – our visual experience will, in Travis’s terms, accept their difference in length as so, and will do this independently of any belief we exercise on the matter. It is not, therefore, the case that autorepresentation functioning as a registration precludes our ability to accept or reject the deliverances of perceptual experience. This conception also does not preclude the contents of representations as being assessable for accuracy in a manner that upholds their independence from the subject’s doxastic state: if the representational content of experience e is simply that p , and this is taken to be so because of how autorepresentation works, then whether this content is accurate or not can still be determined by whether or not p obtains in the mind-independent world.

It may nevertheless be argued that when representation occurs in perception, it *de facto* occurs independently of the subject’s doxastic states, and so Travis’s second premise does go through and autorepresentation is ruled out. Whilst respecting the notion that beliefs cannot influence perceptual experiences without limit, there is still some evidence to the contrary. Returning to the Müller-Lyer example, one proposed explanation of our experience of the lines as different in length is that the hashes on the end trigger our mechanism of depth perception, such that the line with the outward-pointing hashes is interpreted by our visual system as being further away than the one with inward-pointing hashes and, since they are in fact

the same length, these two factors together result in one appearing longer than the other.¹¹¹

In particular, this is because the hashes are thought to resemble common depth-indicative angles as seen in architecture familiar to us. The outward-facing hashes suggest a line that is at the end of a plane furthest from us, whilst the inward-facing hashes suggest a line that is at the end of a plane closest to us. The thought would be that those familiar with these angles and their indications of depth in this way have, in appreciating the typical interplay between depth, distances, lengths, angles and so on in this way, come to learn facts about these properties when they typically come together, and so beliefs with contents involving these facts have become part of – or *penetrated* – their cognitive system.¹¹² As a result of this alleged cognitive penetration, it may be that many of us who experience the lines as being different lengths cannot help imagining the lines as being further away or closer to us due to our familiarity with the relevant angles between the lines and hashes in our highly geometric architecture; consequently, our imagination may combine with the experience of the two lines to yield an experience in which the lines appear to be of different lengths.¹¹³

A point of note about this explanation of the Müller-Lyer illusion is that it is reliant on familiarity with certain angles in architecture for its persuasiveness. Consequently, one would expect that if this were indeed the correct explanation of the illusion, the experience of subjects less familiar with these angles in their surrounding architectural environment would differ. This is empirical question, and there has been some study undertaken to assess differences in experience of the Müller-Lyer illusion across subjects from cultures with varying architectural styles. There were findings of the sort envisaged above, whereby subjects from cultures with less angular architecture were reported as being virtually immune to the illusion.¹¹⁴

¹¹¹ Gregory (1997)

¹¹² Macpherson (2012: 34)

¹¹³ Ibid. (56)

¹¹⁴ See McCauley & Henrich (2006).

There are several ways that the findings of studies like this can be called into question – for instance, it seems reliant on subjects’ self-reporting, the reliability of which may be questioned. Despite not conclusively settling the issue, however, I do take the present example to provide at least some empirical evidence against Travis’s contention perceptual representation must, if it occurs, be allorepresentation that occurs independently of a subject’s doxastic states. Between this and Brogaard’s argument considered earlier, I take it to at least be not obvious that Travis’s argument is successful, and therefore that representationalism still merits serious consideration as a front-running philosophical theory of perception. On this outlook, the question of whether perceptual representation is autorepresentation or allorepresentation remains. I offer now one final consideration in favour of the former over the latter before moving on.

Allorepresentation, as Travis conceives of it, will provide us as subjects with a representation – one that we may then accept or reject. In this respect, allorepresentation is like the English sentence ‘pigs swim’, in that the mere sentence itself can be utilised as an assertion of how things are, or not, and in the former case this assertion of how things are can in turn be accepted or rejected.¹¹⁵ The language of English is not saddled with a commitment to ‘pigs swim’ being how things are simply because it can compose that sentence, and Travis thinks that in just the same way our perceptual representation would not be saddled with being committed to representing things as so in every case, but rather only in certain cases of *committed* representation, in which what is being represented is accepted as how things are. As we saw, Travis argues that autorepresentation involves the commitment to what is represented being how things are: “to autorepresent something just is to accept it [as how things are].”¹¹⁶

If this is the distinction between types of representation that Travis wishes to draw, I submit the following consideration in favour of autorepresentation. I presume that Travis does not wish to deny that we may

¹¹⁵ Travis (2004: 61)

¹¹⁶ Ibid. (62)

believe or disbelieve the deliverance of our perceptual experience (however this is taken to work). I may come to believe that I am experiencing a hallucination or illusion, and may therefore reject the information being provided by my senses, like the Müller-Lyer subject who knows that the lines are of equal length. This process must be independent of the perceptual experience itself, by Travis's own lights. Supposing that perceptual representation is allorepresentation, the representation might be accepted as how things are or rejected, and this process must be independent of the aforementioned belief and disbelief about the deliverances of perception.

On the face of it, this simply does not square with our experience at all. Once I experience something, I may accept or reject that things are the particular way that I am experiencing them to be, but this is an exercising of my doxastic capabilities. In terms of the perceptual aspect to my experience, things have already been accepted as being a particular way, and it is this that I can believe to be how things are or not. This alone does not entail autorepresentation, for it may just be committed representation: a central case of allorepresentation that actually does represent things as being a particular way.¹¹⁷

The difference between committed allorepresentation and autorepresentation is not obvious, however one fact about autorepresentation is that it centrally involves the subject representing something to themselves as being so. I take this to be a particularly plausible explanation for what goes on when we perceive *kind properties*. It seems strange to suggest that my perceiving something as being a doorknob – as I can for variously shaped metal objects in a hardware store – is somehow a property intrinsic to the metallic mass, which could just as easily have been fashioned into something else. If perceptual representation of kind properties occurs at all, it is far more plausible that it is autorepresentation. I return to the issue of perceiving kind properties in Chapter 4.

¹¹⁷ Ibid. (61)

I hope in this section to have illustrated that, unlike Charles Travis, we should consider perceptual representation to be a plausible occurrence and, therefore, should consider representationalism as a plausible candidate for our philosophical theory of perception. In considering Travis's arguments, I have also touched on a distinction between autorepresentation and allorepresentation. Against Travis, I argue that autorepresentation is the more plausible of the two, as it squares more intuitively with our lived experience, and because his argument for the opposite conclusion can be challenged on both theoretical and empirical grounds.

1.4: *Misrepresentation*

I turn now to address the representationalist account of non-veridical experience, which must be supplied if the theory is to avoid the arguments from hallucination and illusion (presented in §1.1), according to which perceptual experience cannot give us awareness of mind-independent objects.

Broadly, the representationalist response to such cases hinges on the view's focus upon the intentionality of experience – the directed 'aboutness' of it – a notion made salient in philosophical discourse initially by Franz Brentano.¹¹⁸ A consequence of the view that experience is intentional which is relevant to the present discussion is that experience of objects is not relational in the sense that the objects experienced serve as constituents of the experience.¹¹⁹ Thus, on this view, "experience is not itself what is intentionally present 'in' it."¹²⁰ As a consequence of this, there is no entailment from an experience of an object to any object of the sort existing; "an experience," to quote Husserl, "may be present in consciousness together with its intention, although its object does not exist at all, and is

¹¹⁸ See Huemer (2014)

¹¹⁹ I discuss this sort of view at length in the following chapter.

¹²⁰ Husserl (1900/1901: 84)

perhaps incapable of existence.”¹²¹ This insight distinguishes the representationalist position not only from the relational view to be discussed in the following chapter, but also from sense-datum theory according to which experience is a relation to some non-physical sense-data.¹²²

Since representationalism appeals to the intentional aspect of experience in this way, a fact about perceptual experience that is brought to the fore as a result is that perceptual experience can be accurate or inaccurate: the experience has *accuracy conditions*. According to the view, it is the satisfaction of these accuracy conditions – in turn – which entails a perceptual experience being veridical: that is, an experience is veridical iff its content matches the way the world actually is in the appropriate way.¹²³ To put it another way: a perceptual experience on the representationalist view leaves open the possibility that the world will match the content of that experience, *and* the possibility that the world will fail to match the content that experience; the experience is veridical in the case where the two match, and is non-veridical otherwise. Since the content of the subject’s experience specifies that experience’s accuracy conditions, the difference between veridical and non-veridical experiences is not located in the content, but the state of affairs in the world, which determines whether these accuracy conditions are met or not.¹²⁴

This allows the representationalist a response to the arguments from illusion and hallucination. According to these, non-veridical experiences should be accounted for in the same manner as veridical experiences, and one cannot say of non-veridical experiences that they constitute awareness of mind-independent objects and their properties, so one cannot say this for veridical experiences either. The arguments from illusion and (the earlier articulation of) hallucination, appeal to the Phenomenal Principle (introduced in §1.1), according to which when one undergoes a non-veridical experience

¹²¹ Ibid. (98 – 99)

¹²² Crane & French (2015)

¹²³ Where ‘in the appropriate way’ is designed to rule out so-called veridical hallucinations, where the world and experiential content match by some freak coincidence.

¹²⁴ Schellenberg (2011: 726)

as of an object, there nevertheless exists *something* of which one is aware. We have seen from the above considerations of representationalism's appeal to intentionality that the view can straightforwardly reject the Phenomenal Principle to answer these articulations.

Another articulation of the argument from hallucination simply sticks to the more basic observation that hallucinations do not involve direct awareness of mind-independent objects (at least with regard to the hallucinated elements), and combines this with the contention that whatever story one tells about veridical experience should match that told about non-veridical experience to yield the argument's conclusion. Here, the representationalist accepts the latter contention on the basis that it is read as the *common kind assumption* (CKA):

(CKA): Whatever fundamental kind of mental event occurs when one veridically perceives, the very same kind of event could occur were one hallucinating.¹²⁵

The representationalist accepts this claim and overcomes this version of the argument from hallucination by holding that one's experience of a snow leopard on the path ahead, for example, is accounted for by one representing a snow leopard on the path ahead; this is the story the representationalist tells whether the experience is veridical or hallucinatory. Representationalism can accept this, whilst avoiding the conclusion of the argument from hallucination by distinguishing between two readings of the conclusion:

(C1): Veridical experiences are not *fundamentally cases* of perceptual awareness of mind-independent objects.

(C2): Veridical experiences do not *give us* perceptual awareness of mind-independent objects.¹²⁶

So long as (C2) does not follow, the representationalist can argue, the argument from hallucination can be satisfactorily answered. Consequently,

¹²⁵ Crane & French (2015)

¹²⁶ Ibid.

the representationalist bites the bullet of (C1), and denies that *any* sort of perceptual experience *fundamentally consists in* perceptual awareness of mind-independent objects.¹²⁷

Instead, representationalism holds that what perceptual experiences of *all* types fundamentally consist in is *representation*. As we have seen, this is compatible with both accurate and inaccurate experiences. The view salvages the perceptual awareness of mind-independent objects that we take ourselves to have by appealing to cases where the world matches the representational content of experience. Instances where this fails to occur are deemed *misrepresentations*, and this term broadly covers both illusory cases – where the representational content of one’s experience correctly identifies an object as present, but incorrectly identifies its properties – and hallucinations, where the totality of the apparent object is misrepresented.

Since representational content is a constant across all levels of veridicality and can be identical in both veridical and hallucinatory cases, this provides the representationalist with a rich story regarding the phenomenal character of hallucinations. This is particularly striking in their account of how certain veridical perceptions and hallucinations can be subjectively indistinguishable from one another – the sameness in phenomenology is to be accounted for by the sameness in representational content. I return to this point in Chapter 4.

This chapter has constituted an exegesis of representationalism: the view according to which perceptual experience is fundamentally a matter of a subject representing their environment as being a particular way. The view is founded upon the notion of experience being *intentional*, and utilises *representational content* in order to account for many of the questions facing philosophical theories of perception. There are many different articulations of the view that share this common core. I have discussed several distinctions drawn regarding both the nature of content (Russellian, Fregean, etc.), content determination (internalism versus externalism), as well as types

¹²⁷ Ibid.

of representation (allorepresentation versus autorepresentation), as well as providing some dialectical context-setting for the view, and illustrating how it is thought to deal promisingly with certain problems in the philosophy of perception. Having thereby provided a thoroughgoing exegesis of one of the two theories in the debate I am centrally concerning myself with, I devote chapter that follows to completing the same task for the other.

2: *Relationalism*

The preceding chapter introduced and provided an exegesis of one party of the representationalism versus relationalism debate. With that analysis of features of and motivations for representationalism complete, this chapter will turn to the task of providing the same for the other party to the debate: relationalism. The chapter will have four parts. The first part serves to provide some considerations designed to provide context for what is known as the relational view, or *relationalism* – i.e. the contemporary articulation of the direct, or *naïve realist* position to which representationalism was opposed. This context will illustrate the place in the philosophy of perception dialectic of contemporary relationalist theories. The second part provides a more thoroughgoing exegesis of this sort of view by considering two influential articulations of it, due respectively to John Campbell and Bill Brewer, with some further analysis provided of the relation of direct acquaintance that lies at the heart of each articulation. The penultimate third part discusses the motivations and merits of relationalism in instances of veridical perception. Finally, the fourth part provides an analysis of the ways in which relationalism deals with non-veridical experiences, discussing the varieties of the relationalist position known as *disjunctivism* about perceptual experiences.

2.1: *A dialectical context for relationalism*

The preceding chapter set the context for representationalism as arising in part out of dissatisfaction and concern about some of the less palatable commitments of direct realism. These theories, we saw, tried to argue that perception was fundamentally a matter of direct acquaintance with the object perceived. However, in order to accommodate the possibility of subjectively indistinguishable hallucinations and the reality of illusions, proponents of these views concluded that the object with which the subject

had direct acquaintance was not, in fact, the mind-independent object itself. Upholding such a view led to the advancement of metaphysically questionable positions, and consequently one narrative regarding the advent of representationalism portrays the view as a response to the desire to avoid such commitments. Another narrative, that particularly relevant to intentionalism, focuses on the interplay between the intentional properties and the phenomenal properties of experience. It is this narrative that will also serve to conveniently carve out a place in the philosophy of perception dialectic for the contemporary form of direct realism, which I am calling *relationalism*.¹²⁸

Before I proceed, I must first present an important caveat that relates the scene-setting above to the discussion below. For all that has been said, I do not mean to suggest that the discussion to follow in this section – which relates among other things to phenomenal properties, qualia, and Frank Jackson’s knowledge argument – is to be taken as an accurate depiction of the historical motivations for relationalism.¹²⁹ The considerations I discuss in what follows serve as the focus of discussion instead because they (a) provide a transition between considerations introduced in the previous chapter and new considerations to be addressed in what follows, and (b) throw into sharp relief how the two views – relationalism and representationalism – differ in their explanatory approaches to a central issue in the philosophy of perception. Moreover, they differ in their approach in a manner that will become relevant to discussion in later chapters, and so it is helpful to lay the groundwork for this in the following section.

¹²⁸ I cannot take credit for the name, which I (fallibly) believe to have first appeared in this context in Schellenberg (2011). For my purposes, I appreciate the similarity in form this name has to representationalism.

¹²⁹ If one is interested in the historical context of contemporary relational views of perceptual experience, I recommend Brewer (2011, esp. Ch.1), and the introductory chapter of Campbell’s (2002).

2.1.1: Separatism

One of the key questions that philosophical theories of perception must take a stand on, if we are to view them as adequate, is what the nature of the relationship between the intentional properties of experience (what that experience is about – representational content falls on this side, for example) and the phenomenal properties of experience (the phenomenal character – what the experience is like for the subject). One response to this issue comes in the form of the following thesis:

Separatism: The phenomenal and intentional properties of experience are independent of one another, and are entirely separable.

The intentionalist positions discussed in the preceding chapter, both strong and weak intentionalism, clearly reject this separatist thesis. They hold that a relation of supervenience (weak intentionalism), if not of identity (strong intentionalism), holds between the two sorts of properties. One might fairly hold that this rejection of separatism is a positive feature of such views. Introspecting upon one's own experience, it seems fair to say that there is at least *some* link between what the experience is about and what it is like for one to undergo.

This does not mean that rejecting intentionalism thereby commits one to separatism, however. Theorists in search of a middle ground are required to establish some sort of link between the intentional and phenomenal properties of experience – one that is non-trivial and goes beyond 'they are both experiential properties', for example – yet also to stop short of suggesting that a perceptual experience's phenomenal character is wholly determined by its representational content. It is possible, for example, to hold the following:

- i) *There are concepts that a subject can only possess if they have perceptual experiences with phenomenal character, and*
- ii) *There are representational contents, constituted by such concepts, which can only be attributed to a subject if that subject has perceptual experiences with phenomenal character, while also maintaining that*

- iii) *We need to invoke the instantiation of non-representational phenomenal properties of experience – i.e. qualia – if we are to account for the phenomenal character of conscious perceptual experience.*¹³⁰

Here, the first two propositions ensure a rejection of separatism by maintaining that representational content is essentially grounded in phenomenal character. Possessing the concepts constitutive of the representational contents of one's experience requires that one must undergo experiences with phenomenal character. The third proposition ensures that the view does not collapse into intentionalism: the latter accounts for phenomenal character in terms of the *representational* properties of experience, yet this view does not.

Instead, phenomenal character is to be accounted for by invoking the notion of *qualia*. Qualia (plural of *quale*) are properties of one's phenomenal experience which are *non-representational*. Strong intentionalists, who are committed to the phenomenal character of an experience being identical to that experience's representational content, will of course reject the notion that such properties exist. Their view simply leaves no room for such experiential properties as qualia. This is not to say that there are not independent reasons for rejecting the existence of qualia, however. In the next section, I will discuss such reasons in more detail, with a view to elucidating the relationalist position to be discussed in the remainder of the chapter.

2.1.2: Qualia & Mary

As we have seen, the qualia theorist can reject the strong claim of separatism – that intentional and phenomenal properties are totally separable from one another – by maintaining that the representational content of a given perceptual experience is explained by its phenomenal character, and then by accounting for the latter in terms of non-

¹³⁰ Soteriou (2016: 92)

representational, phenomenal properties (i.e. qualia) that the experience also has. Whilst it does allow one to do so, this view has not enjoyed universal support from those who wish to reject separatism. In addition to intentionalists who are qualia sceptics because their view leaves no room for these experiential properties, there are those who reject the qualia theorist's position on independent grounds. In this section, I shall discuss briefly reasons both for and against this qualia view, before progressing to the next section to discuss where relationalism fits into the picture of both qualia and the broader (anti-)separatism question.

Support for perceptual experiences possessing qualia as part of their phenomenal properties can be found in the work of self-identified “qualia freak” Frank Jackson.¹³¹ Broadly speaking, Jackson appears to find intuitive the idea that knowing all of the *physical* facts related to, say, feeling pain or experiencing jealousy simply would not confer upon one all that there is to know about the hurt of that pain or the pangs of that jealousy. He recognises, however, that this intuition-based rejection of physicalism is not intuitively obvious to all, and so undertakes the task of providing a valid argument such that even the most enthusiastic physicalists must assent to its soundness.¹³²

In his familiar (1982) example, Jackson asks us to consider Mary:

*[. . .] a brilliant scientist who is, for whatever reason, forced to investigate the world from a black and white room via a black and white television monitor. She specialises in the neurophysiology of vision and acquires, let us suppose, all the physical information there is to obtain about what goes on when we see ripe tomatoes, or the sky, and use terms like 'red', 'blue', and so on.*¹³³

If the physicalist (and so, with them, the intentionalist and qualia sceptic) is correct, then knowing these facts constitutes knowing all there is to know about the phenomenal character of the experience of seeing the colour red, for example. If so, then in coming to learn all

¹³¹ (1982: 127)

¹³² *Ibid.* (128)

¹³³ (130)

about these concepts and physical facts, Mary should be in precisely the same epistemic state with respect to the phenomenal properties of an experience of red as a subject who has undergone such an experience, despite being deprived of phenomenal colour experience.

The natural question to ask, says Jackson, is consequently whether – upon being released from the black and white room and seeing a red thing for the first time – Mary *learns* anything.¹³⁴ In his Knowledge Argument, Jackson contends that it is obvious that Mary will, upon experiencing colour for the first time, “learn something about the world and our visual experience of it,” and so despite knowing all that there was to know in terms of the physical facts of colour experiences, Mary did not know all that there was to know about colour experiences in general.¹³⁵ The qualia theorist stands ready to accommodate this intuition. They can claim that what has happened here is that – despite being in possession of all of the physical facts, and thereby representational properties, of experiencing colour – Mary has for the first time become acquainted with some *non-representational* phenomenal properties, only available to those undergoing colour experience. She becomes acquainted, in other words, with the *qualia* of her colour experience.

So far, so good for the qualia freaks. Let us now assume, however, that Jackson’s picture is correct. This allows for the following adaptation: if it is known that colour perceivers become acquainted with qualia upon perceiving colours, then it is plausible that Mary could come to add this fact to her body of knowledge about colour perception in an entirely theoretical manner (that is: despite her own lack of qualia acquaintance).¹³⁶ Assuming that Jackson’s thought experiment is correct and that this is a plausible assumption to add to it, then upon seeing colours for the first time, Mary gains a certain sort of knowledge

¹³⁴ Ibid.

¹³⁵ Ibid.

¹³⁶ Soteriou (2016: 91)

regarding qualia. Yet she already possessed *theoretical* knowledge about qualia, so what knowledge is it that she has gained? A plausible answer would be that she has gained knowledge about the *intrinsic nature* of the qualia by becoming directly acquainted with it.

The problem with this position is that what experience is like for a subject simply does not cohere with it: if new knowledge of qualia is gained via the subject's experience, then they should be able to introspectively attend to their new knowledge – to the qualia; what the subject discovers, however, is in fact just knowledge regarding properties of physical objects themselves.¹³⁷ If this is correct, it would seem to suggest that any knowledge gained by acquaintance with qualia is actually not introspectively (and thereby mentally) accessible to the subject. Combining this with Jackson's claim that qualia are *epiphenomenal* (that is to say: they are causally impotent with regards to the physical world) and one can argue based on the conjunction of these claims that it is not obvious what explanatory role qualia play, beyond enabling an *ad hoc* denial of separatism.

This is far from a complete account of the debate surrounding the (a) existence of, (b) nature of, or (c) plausibility of (the epiphenomenalism of) qualia.¹³⁸ It is, I hope, nevertheless enough foundation to progress to the next section which situates relationalism within the context of these considerations, providing sufficient background before moving on to provide a thorough exegesis of that view.

2.1.3: *Mary and Naïve Realism*

Returning to the question of what new sort of knowledge one would acquire were one to be in a similar position to Mary (including

¹³⁷ See Campbell (2005) for arguments specific to colour, and (2002, esp. Ch. 7) for a more general argument about knowledge gained through experience.

¹³⁸ Fully exploring such a debate, however, is not the point of this chapter.

potentially complete theoretical knowledge of qualia), one certainly does not *seem* to be newly acquainted with qualia. Recall that these are supposed to be phenomenal properties of one's experience which are non-representational. By hypothesis, one has complete knowledge of the representational phenomenal properties that such an experience would have, and therefore it cannot be the case that one is newly acquainted with such properties.

Campbell's view proposes a middle ground between these two options by accepting each of them, which together yield acceptance of the claim that a subject in such a situation is not newly acquainted with a property of *experience*. Rather, he holds that this newly gained knowledge takes the form of acquaintance with properties of *objects*.¹³⁹ Such a view enjoys intuitive plausibility for several reasons. For one thing, objects and their properties are what a subject like Mary have been newly introduced to, and so accounting for the knowledge gained by appealing to those objects has a certain elegance in its explanatory simplicity. One encounters a red object for the first time and, in so doing, what one gains is knowledge about the colour red itself – as instantiated by some mind-independent object – as opposed to knowledge about one's own phenomenal experiences.

Aside from this intuitive plausibility, the view in question would achieve the following things. First, it would accommodate the intuition that Mary does in fact learn something new upon leaving the black and white lab and seeing the colour red for the first time. She would learn about the property itself, as a property instantiated by objects rather than by her own phenomenal experience. Second, it sustains a rejection of separatism, upholding the idea that there is some important way in which an experience's intentional and phenomenal properties are connected. Yet it does this in a way that does not collapse into intentionalism: the phenomenal properties of experience are accounted

¹³⁹ Campbell (2002; esp. Ch. 7)

for not wholly in terms of representational content, but rather by objects and their properties in the mind-independent world. Moreover, this is all accomplished without reliance on qualia. Though the above discussion was far from conclusive about whether one ought to be sceptical about qualia or not, it is clear that proponents of a view which makes no appeal to such properties need not concern themselves with such debate, as the objections and criticisms brought against the qualia theorist cannot affect the plausibility of their own view.

Though these points may be thought to count in favour of the view in question, for all that has been said so far it remains unclear that this view will avoid concerns shared by many about naïve realist positions. The accommodation of intuitions offered by this view is very similar to that offered by theories of the latter sort. Consequently, contemporary relationalists have been at pains not only to offer an alternative to representationalism, but also to offer such an alternative that accommodates various intuitions that more primitive naïve realist positions were thought to, whilst not falling prey to the same objections as those earlier views did.

The purpose of this part of the chapter was to provide an appropriate historical and dialectical context in which to situate what I am calling relationalism. This took the form of introducing issues such as the question of separatism or Jackson's Mary, which various theories of perception have weighed in on differently based on their respective commitments. By highlighting the ways that other theories – particularly representational theories, with varying views on qualia – have addressed such issues, I have located the place that relationalism finds itself within the dialectic. Having done so, I turn now to provide a more thoroughgoing analysis of the view itself.

2.2: *Relationalism*

As we saw in the preceding section, John Campbell's relational view accommodates many of the intuitions that earlier incarnations of direct realism did. In order to assess motivations for accepting it above any other theory of perception, its workings shall need to be explained in some depth, however. This is the task of this part of the chapter. As we shall see, there are various features of the view, and two prominent articulations of it. In the first section of this part, I will outline the view as portrayed by Campbell, with some remarks on his motivations and aims, as this runs to the heart of contemporary relationalism. Then, in the second section, I provide an exegesis of a variant of the view, due to Bill Brewer, known as the Object View of perception. This is notable for its distinctive treatment of perceptual appearances, but also provides an illuminating analysis of perceptual knowledge that will be useful for my later purposes.

2.2.1: *The relational view*¹⁴⁰

Insofar as the relational view departs from representational ones that claim phenomenal character is determined by representational content, it is itself an alternative account of the phenomenal character of perceptual experience. The view takes as its point of departure the claim that perceptual presentation cannot be wholly reduced to, or manufactured out of, perceptual content.¹⁴¹ In suggesting the impossibility of such complete reduction or manufacturing, Campbell's relational view is due us an alternative account of the phenomenal character of perceptual experience. Quassim Cassam summarises

¹⁴⁰ Though in subsequent chapters I shall use the term 'relationalism' to refer broadly to the sort of contemporary naïve realist position proposed by both philosophers under discussion, for the remainder of this chapter I shall (following them) use *the relational view* to refer to Campbell's relationalist position, and *the Object View* to refer to Brewer's variant thereof.

¹⁴¹ Where perceptual presentation is taken to broadly cover those aspects – purportedly of the physical, mind-independent world – that are 'given' to a subject via their perceptual experiences.

Campbell's take on phenomenal character into the following Constitutive Thesis:

(CT): The qualitative character of sensory experience is constituted by the qualitative character of the objects and properties of the scene observed.¹⁴²

This suggests that the alternative explanation of phenomenal character is to be given in terms of the mind-independent objects in the physical world. In this regard, the relational view is on track to appeal to direct realists.

Concerns raised in the preceding chapter regarding the apparent phenomenal commonalities between veridical and corresponding non-veridical experiences looked to provide reasons to pause before accepting the direct realist conclusion that what we most directly perceive are mind-independent objects themselves. As a result, Campbell's view will need to illustrate that the same fate does not befall it. The remainder of this section first examines exactly how the relational view conceives of perceptual experience – focusing on the workings of the perceptual relation from which the view derives its name – and then proceeds to discuss the role that phenomenal character is then taken to play, before concluding with a discussion on what this does for our perceptual epistemology.

2.2.1.1: The direct acquaintance relation

The name relationalism (or, even more particularly, the relational view) suggests an elevated importance of some relation. The relation in question is the perceptual relation, and to fully appreciate the relational view, one must be aware of what this relation obtains between, and of what the relation itself consists in.

¹⁴² Cassam (2014: 136)

An obvious answer to the first of these points is that the perceptual relation obtains between the subject of the perceptual experience and the object they are perceiving. This is along the lines that rudimentary direct realist positions tended towards. Such views run into problems not only with respect to the arguments from illusion and hallucination, but also from considerations regarding perceiving the same object from different angles (the thought being that the phenomenal character is different from different perspectives, and a simplistic relation between subject and object will be insufficient for explaining this fact). The relation being proposed by the relational view, then, will need to be more complex.

In accordance with this, Campbell contends that it is insufficient simply to hold of perceptual experience *that* subjects stand in some relation with mind-independent objects and properties. Instead, an adequate articulation of the perceptual relation must also take into consideration *how* such subjects stand in that relation. That is to say: an account must be given of how exactly a perceptual experience relates its subject to the worldly objects (and their properties) of which it is a perceptual experience.¹⁴³

The means of factoring how some subject stands in relation to an object into that relation itself, by relationalist lights, involves the expanding of the relation:

*We should think of consciousness of the object not as a two-place relation between a person and an object, but as a three-place relation between a person, a standpoint, and an object. You always experience an object from a standpoint. And you can experience one and the same object from different standpoints.*¹⁴⁴

So, we have the addition of a third relatum. In addition to the subject and the object perceived, we have the notion of a *standpoint*. This is meant to capture the perspectival nature of perceptual experience and thus encompasses the relative position of what is perceived to the perceiver, the sense modality via which they are perceiving, and various other factors that

¹⁴³ Campbell (2014: 51)

¹⁴⁴ Campbell (2011b: 657)

might impact the precise phenomenal experience the subject may have (e.g. lighting conditions in visual perception, or temperature in gustatory perception). In short: the standpoint is as it sounds, and is the perspective from which the subject is perceiving, involving all features of this that work together to influence the exact nature of their perceptual experience.

The addition of this third relatum, intuitively common to all perceptual experiences, accomplishes much of what is required of naïve realist positions such as Campbell's. For one thing, it upholds that perception is a matter of being related to mind-independent objects. The view only differs from this exact statement by emphasising that perception is a matter of being related to mind-independent objects *from a particular standpoint*. In addition, aforementioned concerns regarding how direct realism accounts for the same object being perceived from different angles can now be accommodated, as the relational view can simply invoke the third relatum and say that in such cases the subject and object are held fixed, and perceptual differences simply mirror the difference in standpoint.

As to the nature of the relation itself, it is taken to be a relation of *direct acquaintance* between the subject and object of experience.¹⁴⁵ This at once captures the unmediated nature of the relation between subject and object (from a particular standpoint), and also speaks to the epistemological aspect of perception. Our perceptual experiences *allow* us to know things about our respective environments; they afford us a certain awareness of what is going on around us. In virtue of the perceptual relation being one of direct acquaintance, the thought is that “[w]e have to regard experience of the object as reaching all the way to the object itself, and thereby providing us with the conception of that categorical object.”¹⁴⁶ I return to the business of acquiring knowledge of the *intrinsic character* of object properties below, but presently I continue this part of the discussion with some considerations

¹⁴⁵ Campbell (2002: 115)

¹⁴⁶ *Ibid.* (145)

on what constraints we might be able to place on the *structure* of this direct acquaintance relation so as to further understand its nature.

There are several ways an analysis of the structure of this direct acquaintance relation might proceed. One way may be to focus on its *limits*, as these may offer some insight into the appropriate restrictions to place on our view of its metaphysical structure. This would involve a *contrastive* analysis between what counts as standing in this relation and what does not, in the hope that this would illuminate appropriate delimitations and give us a better understanding of what the relation is capable of encompassing. Another way would be to abstract away from individual sensory modalities, and conduct a *comparative* analysis of what is common to apparent instances of the relation across sensory modalities. The thought, if direct acquaintance is a relation that can obtain in all sensory modalities, is that this would isolate features common to the perceptual relation across these modalities, without which the relation would not obtain.

It is worth pointing out before proceeding that both analyses will proceed via introspection of perceptual experience – a method which has shown itself to be a very poor means of settling the relationalism versus representationalism debate.¹⁴⁷ The goal here is nothing as lofty, however. Instead, the focus will be on what seems introspectively accessible in perceptual experience, as well as what does not – the latter being, itself, a finding that might arguably provide information on the nature of the direct acquaintance relation via establishing its limits. Essentially, the thought is that what we can become aware of via introspection of our perceptual experiences may not be exhausted by *what* we perceive, but might also include the *way in which* we perceive it.¹⁴⁸

One aspect of perceptual experience that we might think is localised to visual perception (at least in some significant way) is that we seem aware of *spatial* relations in our immediate environment and, more broadly, of space

¹⁴⁷ See Spener (2003).

¹⁴⁸ Soteriou (2013: 113)

itself. This is important because it seems as though the space we are aware of need not be occupied by anything in order for us to have this awareness. Seeing a distant object *as* distant seems to involve being aware of a greater deal of space between us and the object than the space we are aware of between us and a nearby object. It seems equally clear, however, that the space we are aware of does not constitute all the space that there is (this is true even within our vicinity: we are usually not visually aware, for instance, of the space directly behind us) – we are aware only of a certain sub-region of it. One may go a step further than this, arguing that the space we experience is space we experience *as being* such a sub-region.¹⁴⁹

If this is correct, then it makes sense to say that there are outer limits to that which we spatially perceive. This seems to correlate well with our visual experience when subjected to introspection: it seems to us as though our visual field consists in certain a more central ‘front and centre’ region, with peripheral vision all sides constituting an outer boundary to our visual sensory field. Considering this boundary, we are not inclined to attribute to it the property of being physical – it does not seem to exist out there in the mind-independent world in the same way as the objects that we perceive. Consequently, it is strictly speaking more plausible consider this a limit of our sensory experience *itself*, rather than a limit of what it is we are experiencing.¹⁵⁰

The first structural aspect of the direct acquaintance relation to be identified, then, is this notion of our *spatial sensory field* in vision. This notion can be thought of as relatively invariant with regards to its structure – it obtains, in some form, irrespective of what object(s) are being perceived. It may be thought of, therefore, as explaining why perceiving spatial locations of objects seems to us to be perceiving some sub region of space in which those objects are located.¹⁵¹ This in turn allows for a neat account of the relationalist claim that the phenomenal character of experience is

¹⁴⁹ Ibid. (118)

¹⁵⁰ Ibid.

¹⁵¹ Ibid. (119)

literally constituted by the objects themselves. The thought is that we have this relatively invariant spatial sensory field, and the objects that fall within it are the objects of which we are visually aware, and are therefore those which can constitute the phenomenal character of that experience.

The areas of our spatial sensory field that are not occupied by objects, meanwhile, still constitute space of which we are aware, yet we are aware of them as empty space in which objects *could* be seen. This story provides us with a way of accounting for the empty space we are aware of, whilst avoiding the need to ascribe to it something like a negative representational content.¹⁵² I take this to be a positive outcome, as the nature of such negative representational contents seems hard to pin down, and would require at least one additional layer of analysis.

If this spatial sensory field were exclusive to vision (I do not argue that it is) then it should, in this respect, to be thought of as isolating a structural aspect of *that* sense modality, as contrasted with others. In general, that we might arrive via contrastive means at an understanding of a structural aspect of particularly visual experience is arguably unsurprising. Not only are humans predominantly visual animals (in the sense that we are mostly reliant on vision to provide us with perceptual knowledge of the world), but the overwhelming majority of literature on perception reflects this, and mostly consists of research into visual perception, at least in cases where it is about one specific sensory modality.

Let us now consider cases that are definitively *not* localised to one specific sensory modality. In accord with the method discussed above, this will involve proceeding via a comparative analysis (across sensory modalities) to try to identify similar limits to our sensory awareness of our surroundings. We saw from our focus on vision specifically that we have a spatial sensory field that constitutes a structural feature of our perceptual awareness by functioning as a limit upon it. It certainly seems true to say that our spatial

¹⁵² Ibid. (117 – 118)

awareness is *primarily* visual.¹⁵³ In identifying limits on our sensory awareness, however, one approach we may try is to consider whether there are any structural features that are equally apparent to *all* sensory modalities.

In this respect, and since we have already identified a spatial sensory field, it makes to consider the notion of a *temporal* sensory field, which will function in a similar way to its spatial counterpart, and which can be demonstrated via hearing. The thought is that our auditory perception is such that we are consciously aware of intervals of time in which sound *could* be heard in just the same way as our spatial sensory field is such that it functions as something we are consciously aware of, and in which objects *could* be seen.¹⁵⁴ Continuing the analogy: just as this characterisation of the spatial sensory field allowed for an account of perceiving empty space when objects did not occupy regions of the spatial sensory field, so too does the characterisation of a temporal sensory field allow for an account of perceiving silence when no objects are heard within the interval constituting our temporal sensory field at that time.¹⁵⁵

Just as before, such a characterisation also allows one to avoid appeal to negative representational content. This is particularly useful if one's view of the representational content of veridical experiences is such that the content stands in a causal relation to objects in the world. We are perfectly capable of envisaging an experience of silence that is veridical, yet it seems this cannot be said to be caused by objects heard, but rather by lack of them. The notion of a temporal sensory field provides an alternative explanation according to which an auditory experience of silence is a positive perceptual experience: rather than being acquainted with a negative representational

¹⁵³ Primacy does not entail exclusivity, though. Hearing a familiar sound at a notably quieter volume – a siren, for example – causes us to believe that there is an emergency vehicle far away from us; if it is gradually getting louder, then we may also come to believe that it is approaching us. Similarly, as it goes past us, it may sound as though it is 'to the left' or 'to the right' (see Lee (2006)). Though fair points, all these really establish is that the spatial sensory field is not limited to vision – a point already established by Gareth Evans's (2002) on Molyneux's question. What is not established here is any new *sort* of limit on sensory awareness.

¹⁵⁴ Ibid. (128)

¹⁵⁵ Ibid.

content, or positively acquainted with a lack of sound, one is instead positively acquainted with an *interval of time* – one in which no sound is heard.¹⁵⁶

Further elucidation of the notion comes from considering when we have the relevant contact with intervals of time. Plausibly, for example, one does not have conscious contact with intervals of time, such that one positively perceives either sound or silence, when one is asleep.¹⁵⁷ That the temporal sensory field extends beyond the auditory can be supported by considering visual perceptions of slight movements: the movement of the hour hand on a clock face, for example. Of course, if we look away for an hour or two and then look back at the clock face, we will judge the hour hand to have moved. Importantly, though, we have not perceived its gradual motion in the same way that we can with the second hand. Its movements may be so slight and incremental that any discernible movement falls beyond the limits of our temporal sensory field. That is to say, its perceivable movement does not occur wholly within the interval of time with which we are directly acquainted, and consequently we fail to experience the motion of the hour hand, despite it moving, when we continually stare at it.¹⁵⁸

Finally, it seems as though there is a degree of transparency regarding the conscious awareness of intervals of time being discussed. Introspectively, we are arguably unable to temporally locate our experiences of durations of time as being located at temporal points *extrinsic* to those durations of time; that is to say: there is no distinct point of view from which we seem to have conscious awareness over temporal intervals.¹⁵⁹ This seems to be a point that can be upheld whilst remaining ecumenical on the issue of whether we perceive temporal intervals in the form of instants or of more substantial length.¹⁶⁰

¹⁵⁶ Ibid. (128 – 129)

¹⁵⁷ Ibid. (129)

¹⁵⁸ Ibid. (129 – 130)

¹⁵⁹ Ibid. (131 – 132)

¹⁶⁰ An issue that crops up in the literature on the specious present concerns how its contents – that “house the change and persistence we encounter in our immediate experience” – relate to ordinary objective time

Overall, the above analysis of the direct acquaintance relation has yielded two potential structural features of our perceptual relation to the world around us. The first of these pertained to conscious awareness of spatial properties, and was cashed out in terms of a spatial sensory field; the second pertained to conscious awareness of temporal properties, and was cashed out in terms of a temporal sensory field. Though possessing differences, which were discussed above, both sensory fields illuminated structural limitations on our conscious perceptual awareness, and each provided a satisfying account of more ‘negative’ perceptual experiences such as seeing empty space or hearing silence. This concludes my analysis of the relation of direct acquaintance appealed to by the relational view, and to the Object View variant to be discussed in a later section. I turn now to consider the role that the phenomenal character plays in the relational view.

2.2.1.2: The role of phenomenal character

The phenomenal character of perceptual experience, on the relational view, is actually *constituted* by properties of the mind-independent objects themselves. We saw this in the above statement of the Constitutive Thesis (CT) – whatever qualitative properties the object in question possesses, is it by these properties that the phenomenal character in question is constituted. This direct link between the subjects of perceptual experiences and the mind-independent objects in the world that they are perceiving allows the relationalist to draw appealing epistemological conclusions on the back of phenomenal character – as the direct relation of perceptual acquaintance is inseparably bound up with phenomenal character, according to the relational view. Establishing exactly what epistemological role

(Dainton (2011: 395)). The literature divides between *extensionalists* (who hold that these contents are temporally extended in the fashion they appear to be) and *retentionalists* (who hold that they are *not* temporally extended in this fashion, yet have the appearance of being so). Extensionalist views have been advanced by John Foster (1982, 1991), Ian Phillips (2010), Barry Dainton (2000) and others; retentionalist views are advanced by Edmund Husserl (1991) and Geoffrey Lee (2014).

phenomenal character plays on the relational view is the task of this subsection.

Perhaps most importantly to proponents of the view, phenomenal character as the relational view conceives of it not only allows us to have knowledge about various truths regarding our respective environments, but also plays the vital epistemological role of accounting for how we come to even possess the *concepts* that we have of the objects and properties that we perceive.¹⁶¹ Not only does the relational view hold that phenomenal character provides us with perceptual knowledge, but it also furnishes us with the very tools required in order to possess and articulate such knowledge in the first place. Another way of putting this is that phenomenal character plays the role of *grounding* our knowledge of the world around us.

One consequence of elevating phenomenal character to play this role is that experience of objects becomes more primitive than our ability to think about those objects.¹⁶² It is this primacy that allows the relationalist to contend that, when Mary sees the colour red for the first time, she is not newly acquainted with qualia, but rather with what the colour red is like, as well as what it is like to see the colour red. According to the relational view, she learns what it is like to see the colour red *because* she becomes acquainted with the colour red itself and learns what it itself is like.¹⁶³ The relationalist will view this as a feature of their view rather than a problem, as they submit that one aspect of the way that we can think of mind-independent objects we perceive is of them *as being* mind-independent.

The role of being that which lets us conceive of mind-independent *as* mind-independent is heavily intertwined with the notion of concept possession. Representationalists hold that this role is played by representational content, taken on several prominent articulations of the view to itself be conceptual. The relational view, conversely, holds that this role is played by

¹⁶¹ Campbell (2002: 122)

¹⁶² Ibid.

¹⁶³ Soteriou (2016: 95)

phenomenal character directly, itself accounted for in terms of acquaintance with mind-independent objects from a given standpoint.¹⁶⁴ It bears saying that the relationalist need not be wholly sceptical about representational content. The mere existence of representational content is perfectly compatible with relationalism; what is not compatible with (austere) relationalism is that representational content plays a prominent role in *explaining* perceptual experience.¹⁶⁵ The two views disagree about what exactly explains this grounding role of perception, but it is precisely this grounding role of perception that the relational view accounts for in terms of the phenomenal character of experience itself, in turn constituted by mind-independent objects and their property instantiations.

One final point to note about this role that the relational view ascribes phenomenal character is that it is this that preserves the relationalist's rejection of separatism. The intentional properties of experience are indeed linked to the phenomenal character of the experience – indeed they are explained by it on the relational view – however the phenomenal character is not itself accounted for in terms of representational content, but rather in terms of the mind-independent objects perceived. Consequently, the view constitutes a theory of perception that is distinct from representationalism, yet which rejects separatism nonetheless.

2.2.2: *The Object View variant*

Before moving on to discuss the motivations and points in favour of relationalism, I wish to draw attention to a variant on the relational view, namely: Bill Brewer's Object View (OV). Brewer's view agrees with Campbell's that *that* we view a mind-independent object is insufficient for a satisfactory theory of perception – we also need to bring in *how* we view the object. In

¹⁶⁴ I do not attempt to argue for one side or the other on this point. See Campbell & Cassam (2014) for discussion of the topic.

¹⁶⁵ For clarity: austere relationalism holds that the perceptual acquaintance relation, and *not* representation, is what does the crucial explanatory work in our philosophical theory of perception.

this regard, the claim that the qualitative character of a sensory experience is constituted by the qualitative character of the objects and properties of the scene experienced is not simply an appeal to the properties and characteristics of the mind-independent objects being perceived themselves, but rather how those objects and properties present to the subject S, standing as S does in relation to those objects and properties *in particular viewing conditions* – i.e. the standpoint relatum discussed above.

I do not intend to rehash in this section an exegesis of how (OV) works when it is in complete agreement with Campbell's relational view. Instead, I shall focus on where the two vary. At its core, Brewer's view is a relationalist view as much and as legitimately as Campbell's is; (OV) simply has its own way of cashing out certain points that not all relationalists would wish to assent to, but explaining these is nevertheless prudent in setting up for the discussion that follows. Consequently, it is this that I shall be focusing on in this section.

As alluded to above, the views agree on perception being a three-place relation involving the subject S, some mind-independent object *o*, and a particular standpoint such that S is directly acquainted with *o* from a standpoint, or within particular circumstances of perception *c*. Recall that it is variation in standpoint that accounts for the same object being viewed from multiple angles (possibly, but not necessarily, by different subjects). Intuitively, we would expect the same object to *look* different from different standpoints. One of the key points of Brewer's (OV) is that it offers a means of cashing out perceptual looks or appearances within a relationalist framework. The remainder of this section will be spent explaining how this works and what consequences fall out of it.

Although the conscious relation of acquaintance proposed by (OV) is, according to the chief proponent of the view, unanalysable, what it is for an object *o* to look some way F to a subject S is analysed in terms of what are called *thin looks* and *thick looks*. Brewer contends that “an object *o* thinly looks F iff *o* has, from the viewing conditions in question, appropriate

visually relevant similarities with paradigm exemplars of *F*.”¹⁶⁶ Here, visually relevant similarities are taken to be similarities with respect to various processes “enabling and subserving visual acquaintance”, e.g. similarities in how light reflects and transmits from the objects present in the scene before *S*, as well as in how such stimuli are processed by *S*’s visual system given their evolutionary history, and so on.¹⁶⁷ In other words, visually relevant similarities simply are similarities with respect to those features in the world that are salient to vision. One can imagine a similar account being given of, e.g., *auditorily* relevant similarities, where these would be similarities with respect to those features of the world that facilitate hearing – similarities in how sound waves interact with the acoustics of their immediate surroundings, and how the subject has evolved to interpret this, etc.¹⁶⁸

A great deal of the work in supporting Brewer’s account of looks is done by the paradigm exemplars to which *o* is to have visually relevant similarities. These are taken to be instances of the physical kinds “whose association with the terms for those kinds partially constitutes our understanding of those terms, given our training in the acquisition of the relevant concepts.”¹⁶⁹ The thought here is that in coming to grasp and subsequently correctly apply the relevant concept of *F*, there are certain paradigm instances of *F* that underpin this. For example, suppose I say that a doughnut looks circular to me. According to (OV), the doughnut must bear visually relevant similarities to paradigm exemplars of circular things, from the standpoint from which I am seeing it. This is to say that it must bear visually relevant similarities to those physical manifestations of circular things which played a central role in allowing me to grasp and correctly apply the concept ‘circular’ to objects.

Meanwhile, *o* *thickly* looks *F* iff *o* *thinly* looks *F* and the subject registers the visually relevant similarities to the paradigm exemplars of *F* that this entails,

¹⁶⁶ Brewer (2017: 216)

¹⁶⁷ (ibid.)

¹⁶⁸ I am merely using this example to elucidate the notion of visually relevant similarities. I am not ascribing this to Brewer’s (OV), though I see no reason in principle why the view could not assent to such a position.

¹⁶⁹ Brewer (2011: 104)

where registration can be thought of as potentially involving the deployment of the concept of *F*, though room is left for other forms of registration that do not involve deployment of the concept in question.¹⁷⁰ The thought behind the distinction is that some object *o* may bear visually relevant similarities to paradigm exemplars of *F*, yet subject *S* may fail to realise this. In such cases, we wish to say that the subject still sees the object, they merely do not see it *as being F*. Brewer reserves the notion of *o* thickly looking *F* to *S* for those occasions where *o* bears visually relevant similarities to *F*, and *S* does appreciate and endorse this fact. In this case, it is not just that *S* is in a position to see *o* as being *F*, *S* actually *does* see *o* as being *F*.

This structuring and cashing out of perceptual appearances allows (OV) to provide a pleasingly straightforward yet intuitively plausible story about what happens in cases of illusion. The view maintains the view of illusion that states that in illusory experiences the subject does indeed experience an object, but does so incorrectly. Given (OV)'s account of looks, it becomes fairly straightforward to give an account of a subject experiencing an object as being a certain way when it is actually not that way. This is because, as Brewer observes, it does not follow from perception being taken to be a relation of direct acquaintance with an object *o* that, in instances of *S* being acquainted with *o* which looks *F* to *S*, *o* actually *is F*.¹⁷¹ Using the apparatus described above, (OV) can accommodate the notion that *from a given standpoint* an object *o* can seem to *S* to bear visually relevant similarities to paradigm exemplars to *F* (when *o* is in fact *G*), and this mistaken association can be attributed in turn to features of the standpoint from which *S* is acquainted with *o*. A basic example would be a blue car appearing to be green thanks to the lighting conditions in which it is being seen.

I return to the relationalist's accounts of non-veridical experiences, both illusory and hallucinatory, in subsequent sections. This also concludes my exegesis of the Object View and indeed of relationalism as a whole. The next part of this chapter proceeds to consider motivations for relationalism:

¹⁷⁰ Brewer (2017: 216)

¹⁷¹ (2011: 108)

arguments for why one might be inclined to accept it, and what further positive things fall out of the theory.

2.3: *Motivations in veridical cases*

In this part of the chapter, I set out and consider some arguments that have been thought to support relationalism in its role as a theory of *veridical* experience. What the relationalist ought to say about non-veridical experience has been a matter of considerable debate, and the following part of the chapter is devoted in its entirety to that question.

2.3.1: *The grounding role of experience*

Referring to objects in the world is an exceptionally commonplace activity; that we are generally able to do this is not in question. It is also an uncontroversial assertion that, in some cases, distinct mind-independent objects of the sort to which we might refer can seem so similar that we cannot tell from their qualitative appearance alone which is which. An obvious example of this is the case of identical twins, however I take it to apply in a similar way to mass-produced everyday items, such as a pair of plain white cups. In each of these cases, if we were to cease attending to one of the pair and – unbeknownst to us – the two were switched, we would be none the wiser on the basis of appearances alone. Subjectively, however, it seems to us in each case as though we are referring to one of the pair in particular: we are referring to *that* cup, or to *that* twin.

Concentrating now on visually identical cups, Campbell observes that when we are each having an experience distinct white cups of this sort, I know which cup I am talking about when I say ‘that cup’, and you know which cup you are talking about when you say ‘that cup’.¹⁷² In considering the source of

¹⁷² Campbell (2002: 124)

this knowledge, a natural candidate seems to be our respective experiences of those particular cups.¹⁷³ If experience does play this explanatory role of grounding our capacity for demonstrative reference, claims Campbell, then this serves as a point in favour of relationalism over representationalism.

This contention is supported by considering how representationalism is committed to interpreting the preceding example. Assuming there is in fact a white cup in front of me, this view commits to my perceptual awareness of the white cup in front of me being a matter of representational content. This content will presumably feature the notion of a white cup, together with any further ways in which I am experiencing the cup in question. The problem is that your experience of a qualitatively indistinguishable white cup is, according to representationalism, to be accounted for via precisely the same representational content. In each case, that there is a white cup and it is in front of us will feature in the representational content, but the point is there is no way for the representational content of my experience to be described that doesn't also serve as a way for the representational content of *your* experience to be described. Regardless of how the example is altered in my case, an analogous scenario can be articulated for yours.

The overall point here is that representationalism looks to be committed to the representational content in each of our respective white cup experiences being identical, despite the mind-independent objects in question being distinct. Yet if this is all there is to perceptual experience – if explanation of it is exhausted by its representational content – then it is unclear how experience can pick out a particular object from a set of qualitatively indistinguishable ones. This squares poorly with the earlier thought that it is our experience of objects that gives us the right to use demonstratives such as '*that* cup'. If experience is to play this role of grounding our demonstrative reference, then it is difficult to see how experience itself could be anything other than a matter of being connected to particular mind-independent objects.¹⁷⁴ If we suppose that instead that experience plays this

¹⁷³ Ibid.

¹⁷⁴ Campbell (2011a: 44)

grounding role and is fundamentally a matter of representational content, then our respective white cup experiences are both a matter of us each representing a white a cup as being in front of us. This representation is meant to account for how it is that we know which particular cup we are each talking about.

The problem with this is that the question of justification is then rebounded, and the representationalist must provide an account of what right we have to use representations in this way. The representationalist seems committed to grounding our use of representations (to be able to refer demonstratively) in representations, and beyond simply building this grounding role in as a fundamental aspect of experience, it is not clear how else they might salvage their account into one possessing explanatory value in this matter. If they opt for claiming this as being an aspect of experience, however, then experience cannot have the explanatory role of grounding our capacity for demonstrative reference. Experience of mind-independent objects, on this view, presupposes our ability to think about them, and so simply cannot be used to explain this capacity.¹⁷⁵ On this basis, it has been argued that only relationalism can provide a satisfactory account of our capacity for demonstrative reference.

2.3.2: The problem of determinacy

This line of reasoning proceeds by observing that since things we experience do seem to be a particular way, there must be some mechanism at work that determines which particular way things seem. With this in mind, it is then argued that relationalism, but not representationalism, can make sense of this. The thought, as Brewer puts it, is that “there are plausible features of the content-like aspects of perception that are quite unintelligible on the assumption that the perceptual system simply serves up specific experiential

¹⁷⁵ Campbell (2002: 123)

contents fully formed and without any explanatory dependence upon more basic experiential facts in particular cases.”¹⁷⁶

There are two considerations at play in the opening formulation of this argument for relationalism, expressible via the following theses:

Variety: For each environmental state of affairs C we perceive, there is a multitude M of ways that it may appear to us (where M is constituted by possible ways of appearing $\langle a_1, a_2, \dots, a_n \rangle$).

Constraint: When we have a perceptual experience of C , our experience constrains M in such a way that only one $a_i \in M$ is how C appears to us.

A toy example that illustrates Variety and Constraint is that of the Necker cube, whose wire-frame nature allows it to appear to us as a cube that is oriented in different ways depending on how we attend to it, or as a tilted hexagon with a symmetrical pattern of lines within it.¹⁷⁷ The three different ways that the Necker cube can appear constitutes the M from Variety, whilst the fact that the cube only ever appears to us in one manner at a time – even if we can bring about aspect shifts by altering how we focus on the image before us – is Constraint at work. Non-illusory examples also work: the table in front of me is such that it may appear to me just to be a table, or an arrangement of planks of wood, or of particles (Variety). In general, though, my experience of the table is such that it only ever appears to me to be one of these ways at any given moment (Constraint), as I will only register one of them at a time. Of course, it may also appear to be green, the colour of emeralds, varnished, large, curved, and so on. The different *sorts* of ways things may appear to be can nest with one another, so that different aspects of the table, such as its size, shape and colour, are each subject to Variety and Constraint.

The preceding quote from Brewer is intended to indicate that only relationalism can account for Variety and Constraint, and that this in itself

¹⁷⁶ (2017: 220)

¹⁷⁷ I am grateful to Fiona Macpherson for the third of these out to me.

constitutes an argument for the view. Views according to which our fundamental account of perception revolves around representational content, as Brewer conceives of them, maintain that the two very different conceptions of looks encapsulated by Variety and Constraint are actually compressed into one layer of representational content.¹⁷⁸ This notion of content, being all that the representationalist has to appeal to, must account for both Variety and Constraint. Brewer's (OV), meanwhile, constructs its account of appearances out of *thin* and *thick* looks, as discussed in §2.2. Thin looks are defined as the visual similarities a perceived object *o* has to paradigm exemplars of various properties. Thus, *o* thinly looks *F* if it bears visually relevant similarities to paradigm exemplars of *F*. Since thin looks are, so described, independent of the subject actually registering the visually relevant similarities, there are a multitude of ways a given object may thinly look. Thin looks are thus ideally suited to account for Variety. Thick looks, conversely, are just those similarities that the subject *does* register, and so can arguably account for Constraint via the notion of conceptual registration via which they are defined. The thought is that conceptual registration on the part of the subject makes determinate the appearance of an object to that subject, and this completes the journey from Variety to Constraint that appears to capture the richness of the world and the particularity of our experience of it.¹⁷⁹

The representationalist might reject this argument on the basis of a two-step argument of their own. The first of these steps is just to accept Variety for what it is: an admission that how things actually appear to us to be is only one out of numerous ways that it might appear to us to be, and this fact operates independently of the particularity of how things actually appear. The second step is to adopt the position that our perceptual experience does involve a particular representational content, and that this content in particular stands in a *causal* relation to the object we are perceiving, thereby satisfying Constraint. In adopting this position, the representationalist could

¹⁷⁸ Ibid.

¹⁷⁹ Ibid. (220 – 221)

arguably complete the journey from Variety to Constraint. They would also be articulating their view in such a way that it is founded upon the notion of representational content, whilst making clear the connection between this content and the mind-independent world.¹⁸⁰

This does not leave the representationalist in the clear, however. The appeal to causation employed in this representationalist response to the worry, the relationalist can counter, simply moves the problem one step up. Now, rather than explaining how the representational content of experience is constrained in the manner described above, the representationalist must explain how the relevant causal relation at work in determining a particular content is constrained to subjects having experiences phenomenally like our own.¹⁸¹

Failure to provide this account leaves their view no closer to satisfying Constraint than it was before. If we accept that experience does work in the manner suggested by Variety and Constraint, and that this is something for which our philosophical theory of perception must account, then it is not clear how the representationalist might answer this challenge. For fear of their view collapsing into relationalism, they should not argue for a more primitive perceptual relation to the mind-independent world. It also seems that they should not appeal to representational content to fix the causal relation discussed above. Doing so would create circularity: the causality is already explaining the determination of the representational content, so determinate representational content cannot explain the causality. Abandoning the causality altogether leaves them once again in the position described by Brewer, where their initial means of accounting for Variety and Constraint is wholly unclear. Consequently, considerations of determinacy arguably motivate relationalism.

¹⁸⁰ One of the objections Brewer raises to what he calls the Content View (CV) in his (2011) is that the view leaves the connection between perception and the mind-independent world it is supposed to be informing us about insufficiently clear.

¹⁸¹ Soteriou (2016: 106 – 107)

2.4: *Disjunctivism*

For all that has been said so far of relationalism as a theory of perception, one could be forgiven for thinking that the view still requires an account of non-veridical experiences. The exposition of the variants of the relational theory discussed above proceed, in the most part, by providing accounts of veridical experience, and focus largely on the attractive epistemological story that such accounts provide, contrasting this with arguable shortcomings that come from embracing representationalism. The Object View's analysis of appearances allowed for an account of illusory experience to be derived from the view's account of veridical experience, but so far as it goes, this achieves little beyond gesturing towards some features of veridical perception, and hoping to explain illusion by simply observing that it lacks these features. More importantly, nothing at all has been said about how the relationalist is to account for hallucination. Infamously, no philosophical theory of perception is to be treated as satisfactory unless it can accommodate the complications that these two sorts of non-veridical experiences often throw into the works.

Relationalism must consequently provide an account of non-veridical experiences. Inevitably, this will differ from the one provided by representationalism, since the latter derives its explanation of illusion and hallucination from the notion of representation, which plays no prominent role in a purely relationalist picture. Also providing restriction on the relationalist account of non-veridical experience are the arguments from illusion and hallucination. Each of these arguments concludes that we never have perceptual experience of mind-independent objects - a conclusion to be resisted by representationalists and relationalists alike. The theories differ in exactly how they resist this conclusion, a fact illustrated by which premises

of the argument from illusion and the argument from hallucination they respectively reject.

One contention that is relevant to both of these arguments is that when a subject undergoes a non-veridical experience, there is arguably nevertheless *something* of which they are aware. Representationalists reject this step, as they will hold that “[it] is not generally true that when a representation represents something as being *F*, there has to actually be something which is *F*.”¹⁸² By rejecting this premise, the representationalist undercuts the arguments in question before the arguments use a spreading step to move from the notion that we are not aware of mind-independent objects in non-veridical experience to the conclusion that we are never aware of mind-independent objects at all.

The spreading step in question suggests that the same account must be given of veridical and non-veridical experiences, and it is this step in the argument that relationalists are inclined to question. This approach contrasts with representationalism, in that the latter is often referred to as a ‘common factor’ theory, and suggests that there is some experiential, intentional feature of qualitative experience that is common to veridical perceptions, illusions and hallucinations. The experiences are therefore thought to be of the same fundamental kind, with the difference between them located in how well the world matches up with their experiential content.

The remainder of this section explores the relationalist alternative to this, broadly known as *disjunctivism*. Traced back to Hinton, the view broadly holds that when one has an experience as of *e*, one is *either* having a veridical perception of *e*, *or* is undergoing a non-veridical experience (an illusion or hallucination) in which it merely *seems* that *e* is present, thus denying that the same account must be given of the so-called ‘good’ (veridical) and ‘bad’ (non-veridical) cases.¹⁸³ The exact way in which this

¹⁸² Crane & French (2015)

¹⁸³ Hinton (1967)

denial is articulated can vary: it can be thought of as merely acknowledging epistemological difference, or as a stronger, more metaphysical claim that the two sorts of experience are of a different psychological kind.¹⁸⁴

Whichever mental aspect serves as the vertex of the two disjuncts, a further claim can be made that what is present in the ‘good case’ disjunct and absent in the ‘bad case’ one is fundamental to that mental aspect.¹⁸⁵ A proponent of this sort of view is also of the view that the good and bad cases are of *fundamentally* different kinds. I acknowledge this here, but as much of what follows in subsequent chapters denies that perceptual experience deals in such fundamentals, I shall not pursue it further. In what follows, I discuss the metaphysical and epistemological versions of disjunctivism in more detail, highlighting what each allows the relationalist to say regarding experiences of hallucination and illusion.

2.4.1: Metaphysical Disjunctivism

The most influential way that the disjunctivist denial of a common factor between veridical perception and hallucination has been articulated is what I am calling *metaphysical* disjunctivism. This position holds that the mental events of veridically perceiving and hallucinating are of different psychological kinds, even when they are subjectively indistinguishable. Since perceptual experiences (veridical perceptions, illusions or hallucinations) are such that they are of some psychological kind or other, by asserting that veridical and hallucinatory experiences are of different psychological kinds, the metaphysical disjunctivist guarantees that the two experiences are to be accounted for in different ways. This, in turn, blocks the so-called spreading step of the arguments from illusion and hallucination, according to which veridical and non-veridical experiences are to be accounted for in the same manner. Note that as I am using the term metaphysical disjunctivism here,

¹⁸⁴ Not an exhaustive list of the varieties of disjunctivism, but this does cover the two by far most influential varieties. For further discussion on others, see Soteriou (2016, 2014), Byrne & Logue (2009), and Haddock & Macpherson (2008).

¹⁸⁵ Soteriou (2014)

the view is broadly equivalent to what Byrne & Logue (2009) call *moderate* disjunctivism, as distinct from their articulation of metaphysical disjunctivism, which is such that there is *no* mental event common to veridical perception and hallucination.¹⁸⁶ I take both views to be legitimate and distinct from one another, but do not think that anything that follows hangs on their distinction. Consequently, I maintain my usage, viewing the difference between the moderate and metaphysical (in their sense) versions of the view as a matter of degree.

This version of disjunctivism is not uncontroversial, however, and one significant challenge raised against the view is to provide an explanation for *causally matching hallucinations*. The idea is this: suppose one veridically perceives a glass of water. Once one goes far enough down the causal chain, it makes sense to say that the physical glass of water in one's immediate environment is a cause of one's veridical experience. In considering the most proximate cause of the veridical experience, it makes sense to conclude that one's brain state caused one's experience, however. The thought then runs that it is conceivable that one might have one's brain artificially brought into this state, and that this would lead to a qualitatively indistinguishable hallucination occurring, which had the same proximate cause as the corresponding veridical perception. Because of this latter point, the argument runs, the two experiences should be accounted for in the same fashion, blocking the disjunctivist solution to the argument from hallucination.¹⁸⁷

Many disjunctivists have rejected this argument, however, citing as a reason their commitment to the relationalist's constitutive claim regarding phenomenal character. The suggestion, identified as the Constitutive Thesis (CT) earlier, was that the phenomenal character of a veridical experience is literally constituted (at least partially) by the mind-independent objects perceived. Consequently, it is open to the disjunctivist to argue that the same mental event cannot occur in both a veridical perception and its

¹⁸⁶ See Byrne & Logue (2009).

¹⁸⁷ Crane & French (2015)

causally matching hallucination. Both mental events have phenomenal character, and (CT) entails that the phenomenal character of the veridical experience is constituted by the mind-independent object(s) perceived. By definition, the corresponding, subjectively indistinguishable, hallucination is such that its phenomenal character cannot be accounted for this way, as there is no salient mind-independent object being perceived. Since the presence of a mind-independent object is essentially constitutive of the phenomenal character of veridical experience, then even if one could artificially bring about the same brain state to achieve a causally matching hallucination, the events would not be of the same psychological kind. This is because an essential constituent of the veridical phenomenal character is missing from the hallucinatory case.¹⁸⁸ As a result, even if the disjunctivist allows that the veridical perception and hallucination *are* causally matching, their relationalist commitment to (CT) allows them to resist the claim that mere causal matching entails a sameness in accounts given for the different experiences.¹⁸⁹

This fails to completely overcome the concerns posed by hallucination. The opponent of disjunctivism can modify the principle at play in their causal argument from ‘if two events, e_1 and e_2 have the same proximate cause, then they are to be explained in the same way’, to something like ‘if events e_1 and e_2 are brought about by the same proximate cause *in circumstances that do not differ in any non-causal conditions necessary for the occurrence of an event of kind K*, then they are both of kind K , and are to be explained in the same way.’¹⁹⁰ It is then possible for them to argue that all that is required for a hallucination *as of* E – let this be of kind H – is that a subject is in a specific brain state, and this brain state would also be present in the corresponding veridical perception of E. Consequently, even if one can argue that a veridical perception of E is of kind V due to being constituted by whatever E is, it is

¹⁸⁸ Martin (2004: 56 – 57)

¹⁸⁹ Crane & French (2015)

¹⁹⁰ Nudds (2009: 337)

also of kind *H* since there is no difference in the non-causal circumstances necessary for the occurrence of the brain state in either case.

From here, the argument continues, we run into a screening off objection. The hallucination of *E* has a phenomenal character that is grounded in its being of kind *H*, and if this is so then it is hard to see how the veridical perception of *E* being of kind *V* explains the phenomenal character of that experience better than its being of kind *H*, since being of kind *H* is itself sufficient to account for this.¹⁹¹ As Robinson puts it, “[if] simple indiscriminability is enough to constitute the phenomenology of hallucination, how could it fail to do the same job for a perception indistinguishable from the hallucination?”¹⁹²

The most popular way for the metaphysical disjunctivist to address this issue is to adopt the so-called negative epistemic conception of hallucination. According to this view, most prominently articulated by Michael Martin, all there is to a hallucination being as of a particular scene with a certain phenomenal character is that it is subjectively indistinguishable through introspective reflection from a veridical perception of that scene.¹⁹³ If that is so, Martin observes:

[The] *property of being a veridical perception of a tree never has an explanatory role, since it is never instantiated without the property of being indiscriminable from such a perception being instantiated as well. But if the property of being a veridical perception lacks any explanatory role, then we can no longer show that being indiscriminable from a veridical perception has the explanatory properties which would screen off the property of being a veridical perception.*¹⁹⁴

In this way, the disjunctivist can utilise the apparent fact that the explanatory power of the hallucination being of kind *H* is *parasitic* on the explanatory power of the corresponding veridical perception being of kind

¹⁹¹ Crane & French (2015)

¹⁹² (2013: 326)

¹⁹³ Martin (2004: 69)

¹⁹⁴ *Ibid.*

V.¹⁹⁵ This arguably allows the disjunctivist to overcome this screening off problem.

In spite of this, the negative epistemic account of hallucination is not without fault. It accounts for hallucination in a manner wholly devoid of a positive account of its phenomenal character. When contrasted with the rich story that the representationalist can provide of hallucination, this seems lacklustre: the representationalist can tell an illuminating story of hallucinatory experiences via the appeal to misrepresentation, underpinned by their common factor approach. In conjunction with their rejection of the notion that this entails that veridical perceptions and hallucinations must be accounted for in an identical fashion, their theory can resist the argument from hallucination. As the disjunctivist attempts to circumvent this argument by denying the common factor, this approach is not available to them. Since the latter view is so minimal, negative and singularly focused on a hallucinating subject's epistemic state, despite its claims that the bad and good cases differ on a psychological and metaphysical level, it is arguably a fair criticism that the view fails to account for all that we wish to explain about hallucination.

The notion of indistinguishability being used to account for causally matching hallucinations, for one thing, is arguably incapable of doing all that it is supposed to. We can see this once we clarify the notion of indistinguishability being used. Recall that, according to Martin, *all* there is to the phenomenal character of such hallucinations is that they are indistinguishable from their corresponding veridical perceptions.¹⁹⁶ Since Martin is committed to providing different accounts of phenomenal character for veridical and non-veridical experiences, it follows from this that on his view there is something more to the phenomenal character of veridical experiences. We have seen above that this something more is constituted by the sensible properties of the mind-independent objects perceived: it is because the objects of veridical perception seem to us to have

¹⁹⁵ Crane & French (2015)

¹⁹⁶ (2006: 369)

the properties they do that we recognise the experience as being phenomenologically the experience that it is.

If this is how we recognise veridical experiences phenomenologically, and they are indistinguishable via reflection from causally matching hallucinations, then the disjunctivist is open to the following problem.

Suppose we recognise a veridical experience, v , as being phenomenologically the experience that is because it seems to present property F . It is plausible that any experience subjectively indistinguishable from v – e.g. causally matching hallucination h – is one that we also recognise phenomenologically because it seems to present property F . If so, Martin must make sense of the following: since h is recognised *as indistinguishable* from v in virtue of appearing to present F , it cannot be that this F appearance is explained wholly by indistinguishability, yet his view seems committed to this.¹⁹⁷ The above characterisation of our recognition of the phenomenology of v is a commitment of disjunctivism. Meanwhile, the claim that the phenomenology of indistinguishable experiences is recognised in the same way seems a plausible truth about indistinguishability.

An alternative approach might be to claim that the subject of h 's impression of how their experience phenomenologically appears is the result of an unconscious inference – one which occurs quickly and delivers to them the impression that their experience is just like the one they would be having were they having v . Given that v has the clear phenomenology of appearing F , however, it seems implausible that this impression of similarity would occur without h also striking the subject as appearing F – a positive phenomenal feature of the sort that Martin's view precludes. Consequently, the only path of resistance for Martin is to insist that h 's seeming to be F is not to be analysed in terms of its possession of the intentional of phenomenal object F , but rather that its having that object should be analysed in terms the experience's indistinguishability from v 's having that object.

¹⁹⁷ Robinson (2013: 324 – 325)

Indistinguishability is not yet out of the woods. Susanna Siegel raises the point that for two experiences to be indistinguishable (as defined by Martin) is for a subject to be unable to judge – by introspective reflection alone – that they are different, and argues that this problematises accounting for cognitively unsophisticated hallucinators. Giving the example of a dog, she claims that since the dog lacks the cognitive sophistication to form judgements, the dog *a fortiori* cannot form judgements of difference via introspective reflection or otherwise, and this will have the consequence that *all* the dog's experiences are indistinguishable.¹⁹⁸ This has the implausible consequence that any dog hallucination will be indistinguishable from each and every dog experience.

In responding to this, Martin argues that the notion of indistinguishability being appealed to in his account is *impersonal* – i.e. it is not due to some deficiency in the subject, and holds of things that are indistinguishable *per se*.¹⁹⁹ This allows his view to circumvent concerns regarding cognitively unsophisticated hallucinators, such as dogs. This response also faces problems. One can ask more of a pair of indistinguishable experiences than whether they are impersonally indistinguishable – one can ask in what respect(s) they are indistinguishable.²⁰⁰ Here, a natural answer for *v* and *h* is that they both appear *F*. Yet this positive feature cannot be accommodated by Martin's view, because on that view *no* positive feature of hallucinations aside from their indistinguishability from counterpart veridical perceptions can be accommodated.²⁰¹

It also seems plausible that we would like a way of picking out which state it is that is (impersonally) indistinguishable from the corresponding veridical perception in an acceptable manner that does not beg the question.²⁰² Here, Martin's view arguably faces further problems because the experience in question cannot be identified by any robust positive trait, for on Martin's

¹⁹⁸ (2008: 210)

¹⁹⁹ Martin (2006: 381)

²⁰⁰ Robinson (2013: 326)

²⁰¹ *Ibid.*

²⁰² Siegel (2008: 212)

view it has none. Similarly, the notion of impersonal indistinguishability cannot help, for “we are trying to explain what it is for a state of the dog’s to have the [indistinguishability] property in the first place, so we cannot get a fix on which state we are talking about by appeal to its being the one that has that property.”²⁰³

I return to the prospects of Martin’s view for explaining all that we want a philosophical theory of perception to explain about hallucination in Chapter 4. I turn now to consider an alternative metaphysical disjunctivist account of hallucination. William Fish discusses a disjunctivist position that is *eliminativist* about phenomenal character: the view maintains that all that we wish to explain about hallucinations can be explained via the beliefs the hallucination brings about in the subject, and these will include the belief that they are having an experience with phenomenal character.²⁰⁴ In short, the cognitive effects of a hallucination account for its apparent phenomenology.

If one accounts for things in this way, then one can maintain that hallucinations in fact lack phenomenal character altogether, whilst explaining hallucinators’ beliefs to the contrary. This arguably allows the disjunctivist to explain the qualitative identity of hallucinations and their corresponding veridical perceptions in such a way that the disjunctivist commitment to there being no shared phenomenal character between the two sorts of experience. After all, if hallucinations lack phenomenal character altogether then they cannot possibly share any with their veridical counterparts, regardless of what particular phenomenal character is possessed by the latter.

Fish’s disjunctivism is committed to the controversial view that the subject is fallible with respect to whether their experiences have phenomenal character – a consequence that brings with it new complications regarding the seems/is distinction. One can be forgiven, for instance, for accepting that

²⁰³ Fish (2009: 92)

²⁰⁴ Ibid. (98)

if my experience *seems* to me to have a phenomenal character, then that is what that experience is like for me, and the experience therefore has phenomenal character by definition. Fish's position is committed to the falsity of this, and therefore owes us an account of our fallibility in this respect.

A more serious issue arises when we consider different sorts of hallucinations. It is fair to say that most philosophical accounts of hallucination are founded upon considerations of what are might be called *philosophers' hallucinations*: mundane hallucinations that are qualitatively indistinguishable from a correspondingly mundane veridical perception. More properly, however, a philosophical conception of hallucination should additionally encapsulate experiences resulting from taking hallucinogenic substances, and experiences which are such that the world could never be as the experience presents it to be, because what is hallucinated could never possess a veridical counterpart.²⁰⁵

According to Fish's view, hallucinations lack phenomenal character. The phenomenal character that they seem to have is simply the result of the hallucination's cognitive effects, and the resulting false belief generated can explain all further behaviour based on the hallucinatory experience. In the case of philosophers' hallucinations, this can be further cashed out in terms of the corresponding veridical experience: subjects of such hallucinations mistakenly believe that they are having an experience with the precise phenomenal character that the corresponding veridical perception would have.

The problem as I see it with Fish's view is that it is not applicable to the other sorts of hallucinations just mentioned. If indeed it is possible to hallucinate impossible figures, or if the hallucinatory visuals resulting from ingestion of hallucinogenic substances do not correspond to a way the world

²⁰⁵ Macpherson (2013: 1); experiences of the latter kind would be hallucinations of the sorts of impossible figures that Escher paintings, for instance, brought about illusions of. See Macpherson (2010) for discussion of hallucinations of this sort.

could possibly be, then it is difficult to see how to accommodate the fact that such experiences *seem* to have phenomenal character. This is because the hallucinations in question do not even seem to be veridical. Fish is committed the view that these hallucinations do not actually possess phenomenal character. Yet if this is to be explained via cognitive effects of the hallucination, and by false beliefs formed, the question arises as to what this false belief is based on? It cannot be on any corresponding veridical perception in these cases for, as was just established, they lack such counterpart veridical experiences.

There are two options I can see for Fish. The first is to commit to a sort of error theory across hallucinatory experiences, providing two different accounts - one for philosophers' hallucinations, one for other hallucinations - which both involve us being wrong about our non-veridical experiences. The reason two accounts are needed is that the earlier point stands: why hallucinations with *no* counterpart veridical experiences seem to have phenomenal character, yet do not, cannot be explained by his account described above. This avoids the issue by separating the two sorts of hallucinations, accounting for each slightly differently, yet maintaining eliminativism about the phenomenal character of hallucinations under the common banner of an error theory.

The second option is to accept that certain hallucinations that seem to have phenomenal character *do* have phenomenal character, and to restrict his eliminativism about the phenomenal character of hallucinations to philosophers' hallucinations, building this hard distinction in to the comparative phenomenology of philosophers' hallucinations and other types. This avoids the issue by conceding that different sorts of hallucinations have a different sort of phenomenology, and restricting eliminativist claims about phenomenal character to philosophers' hallucinations alone.

If Macpherson is right, however, and the overall philosophical conception of hallucination must be inclusive of both sorts of hallucination, then either

approach will lead to the overall account of the philosophical conception of hallucination in question being disunified. The first approach requires two separate accounts in order to construct a single error theory; the second will involve a disunified story of the phenomenology of hallucinations. If we end up with a disunified account of hallucination, then it is unclear what commonalities they possess that merit the two being presented under the banner of 'hallucination'. Insistence that the two are indeed of a common psychological type - hallucination - will not help. For this leads (of all things) to a screening off objection.

This is easiest to see in considering the second approach, according to which some hallucinations can have phenomenal character. If having phenomenal character is sufficient to account for the phenomenology of hallucinations which are not philosophers' hallucinations, then it is hard to see why the relation to veridical counterpart experiences with phenomenal character must step in to play this role only in the case of philosophers' hallucinations in a manner that isn't *ad hoc*. In terms of the first approach, a similar argument could plausibly be deployed by substituting phenomenal character with whatever explains the error theory for hallucinations that are not philosophers' hallucinations.

In order to avoid this screening off objection, the insistence that the two sorts of hallucination are of a common psychological type must be dropped. Yet once this has happened, there are parallels to metaphysical disjunctivism that are difficult to ignore. Recall that one of the base principles of the position was that since veridical perceptions and hallucinations are of different psychological type, they must not be accounted for in the same way. That is to say that they are not to be thought of as the same sort of experience as one another. It is not clear, if the above analysis is correct, why we must adopt this approach for the disunity of types of experience (veridical perception and hallucination) and not for the disunity of hallucinatory experience.

The above is an objection that only applies to Fish's account of hallucination, because only that view involves eliminativism about phenomenal character. To that end, Martin's metaphysical disjunctivism is unaffected by such considerations, because his view permits hallucinations to have phenomenal character across the board, but the story told about the phenomenal character of philosophers' hallucinations is very minimal.

Before proceeding to discuss a different variety of disjunctivism in the next section, it remains to be addressed how the metaphysical disjunctivist ought to account for illusory experiences. As described above, metaphysical disjunctivism divides experience up into two kinds: veridical and illusory. The question naturally arises: on which side of this distinction do illusions fall? Adopting first-letter abbreviations, answers fall mainly into two camps: VI vs H, and V vs IH - where the former groups illusions in with veridical perceptions as being radically different to hallucinations, and the latter groups illusions in with hallucinations as being radically different to veridical perceptions.²⁰⁶ In virtue of illusions failing to be fully veridical, Logue notes that several relationalists adopting metaphysical disjunctivism accept V vs IH without argument.²⁰⁷

As there is nevertheless *some* veridical element to illusions, however, this as it stands seems unsatisfying. For one thing, it leaves open the question of how illusions and hallucinations differ. That being said, it can also be argued that the V vs IH / VI vs H distinction is too crude. Disjunctivism is in essence a negative thesis (that hallucination and veridical experience lack a common nature) that was constructed in order to support a positive one (that veridical perception is fundamentally a matter of the subject bearing a conscious perceptual relation to objects in the world); it does not follow from this that there should be a unified account of all the ways an

²⁰⁶ Logue (2015: 210)

²⁰⁷ Ibid.; Logue correctly identifies Martin (2006: 360) and Campbell (2002: 117) as doing this.

experience may not be veridical: there is no restriction on the number of ways in which things may go wrong in our experience.²⁰⁸

Nevertheless, it is worth examining what disjunctivists are committed to, because of their accounts of hallucination, in terms of accounting for illusion. The view of hallucination offered by Fish, for example, *cannot* be in the V vs IH camp in the simplest sense. This is because his view claims that hallucinations lack phenomenal character altogether. If illusions are like hallucinations in this respect, then it is extremely difficult to see how illusions may be maintained as being partially veridical, as this would seem to suggest at least some phenomenal character to the experience.²⁰⁹

Moreover, if there is experiential continuity between veridical perception and illusion – as seems plausible – then much of our experience is somewhat illusory; placed in conjunction with Fish’s account of hallucination, this implies that we only occasionally enjoy experiences with phenomenal character, and erroneously believe that we do in the vast majority of cases.²¹⁰

In dialectical terms, it is not the case that in virtue of being committed to not endorsing V vs IH, Fish is committed to endorsing VI vs H. Upon closer inspection, one can find ways to make room in logical space for views that are not wholly on one side or the other of that distinction. As Fish puts it:

*The fact that illusions all involve some features of an object being seen correctly and other features being seen wrongly [. . .] should not blind us to the fact that there are significant differences between the sorts of scenarios that fall under the broad banner of illusion.*²¹¹

This hints the broad strategy that Fish’s own view employs: observe that multiple sorts of experiential state are each legitimately classed as illusory, highlight the reasons for difference in a finer-grained classification of these

²⁰⁸ Soteriou (2016: 183)

²⁰⁹ Fish (2009: 146)

²¹⁰ Ibid.

²¹¹ (2009: 147)

types of illusions, and offer different explanations of these types of experiences.

Fish identifies three types of illusions: physical illusions, optical illusions, and cognitive illusions. *Physical* illusions are predictable, intersubjective illusions which occur solely because of the way the mind-independent world is. The example of a straight stick appearing bent in water is given, with the observation that not only can this type of illusion be wholly accounted for without any appeal to what happens after light hits the subject's retinas, but it also has the property of being able to be photographed.²¹² In accounting for these illusions, which are wholly explicable by appeal to worldly states of affairs, Fish concludes that they are actually a special case of veridical perception.²¹³ In particular, they are a case of veridical perception that "involves the successful perception of somewhat unusual facts."²¹⁴ These are taken to be facts which preclude the successful perception of other facts, but are nonetheless themselves veridically perceived.

With regards to the stick in water, this amounts to the water distorting the way in which we see the stick, thus precluding us from perceiving the fact of the stick's straightness, yet enabling us to see how the stick appears when viewed in glass of water from a certain angle.²¹⁵ Consequently, physical illusions can be accounted for in the following general way: in virtue of veridically perceiving some facts, we are precluded from perceiving others. On the basis of veridically perceiving these unusual facts, we might be prone to formulate false beliefs, however explaining this requires the resources of the next type of illusion to be discussed.²¹⁶

Cognitive illusions must be explained by appealing to the particular contribution made by the subject to the misleading appearance, and may be less intersubjective and predictable, to varying degrees, than the other types

²¹² Ibid. (148)

²¹³ Ibid (151)

²¹⁴ Ibid. (171)

²¹⁵ Ibid. (165)

²¹⁶ Ibid.

of illusions.²¹⁷ Fish provides an account of these that is similar to his account of hallucinations, inasmuch as false beliefs are formed that account for the subject's consequent behaviours. Considering as an example a subject with ophidiophobia (fear of snakes) taking a coiled rope to be a snake, Fish argues that the ophidiophobia and accompanying heightened anxiety cause the subject to misdeploy their recognition capacity, leading to the mistaken belief that the coiled rope is a snake.²¹⁸ Once the subject becomes aware of their error, the illusory experience does not persist - in the example given, the ophidiophobe stops illusorily experiencing a snake upon becoming aware that what is there is actually a rope.

Cognitive illusions differ from hallucinations, however, in that they are still causally connected to the mind-independent world. In virtue of this fact, even with an explanation that is wholly rooted in the role of the subject of the experience, there is some constraint upon how erroneous the illusory aspect of the experience can be.²¹⁹ The ophidiophobic subject's condition, together with their viewing of a coiled rope of a greenish-brown colour in certain lighting conditions accounts well for their experience as of a snake. Conversely, it accounts poorly for an experience as of a chessboard. In the latter case, we would be inclined to conclude that the subject was hallucinating, as intuitively there is no such limitation imposed by the mind-independent world onto how erroneous a hallucination can be. I return to this point for different reasons in Chapter 4.

Optical illusions are predictable and intersubjective, as physical illusions were, however they are not explicable without some reference to processing carried out by the subject resulting in the misleading appearance.²²⁰ In this respect, they are like cognitive illusions, in that the subject forms false beliefs about what they are seeing, and these beliefs account for further behaviour the subject exhibits.²²¹ Examples are numerous and familiar in the

²¹⁷ Ibid. (149)

²¹⁸ Ibid. (169)

²¹⁹ Ibid. (170 – 171)

²²⁰ Ibid. (148 – 149)

²²¹ Ibid. (174)

literature, such as the Müller-Lyer illusion and the Adelson chessboard. They differ from cognitive illusions, however, in the sense that the illusory experience can persist after the subject's error has been pointed out to them: once one is aware that the two longest lines in the Müller-Lyer illusion are of the same length, one nevertheless can continue to experience one of them as being longer than the other.²²²

Optical illusions also differ from both physical illusions and cognitive illusions in the following sense. The illusory aspects of physical and cognitive illusions are accounted for wholly by appeal to the mind-independent world and the subject respectively; in contrast, optical illusions will appeal to both. The appeal to the role of the subject is described above. According to Fish, appeal to the worldly aspect will vary in specifics according to the particular optical illusion, but will be to "the relevant illusion-inducing features of the perceived scene."²²³ Returning again to our Müller-Lyer lines, if the subject is presented with only the two longest lines, they will correctly report that they are of the same length. It is when the inward and outward-facing hashes are added to the ends of these lines that the illusory experience occurs. In this case, then, the worldly features appealed to are the hashes, and the processing that these lead the subject to perform constitutes the cognitive element of the explanation of the illusion.

In general, since it divides the broader class of illusory experiences up in this way, Fish takes his view to operate at "the level of individual facts that the subject is acquainted with in experience, rather than at the level of whole experiences."²²⁴ Observing this division into veridical perception of unusual facts and false beliefs about worldly facts on the part of the subject (or a combination of these two), the view itself has a disjunctive element to it. On this view, a subject seemingly acquainted with a worldly fact *either* is acquainted with that fact *or* is mistakenly believing that they are acquainted

²²² Ibid.

²²³ Ibid. (176)

²²⁴ Ibid. (178)

with that fact.²²⁵ In summing up what these two claims mean for the view of illusion, Fish concludes:

*Therefore, at the level of individual phenomenal properties, there is no class of illusion; there are only two ways of seeming to see a particular feature: either veridically or non-veridically. The key to the picture, however, comes with the recognition that any normal visual experience will seem to the subject to present not just one, but an array of worldly facts. At the level of whole experiences, illusions occur when some features are seen veridically and others non-veridically.*²²⁶

As a result of this, Fish's view as a whole can take classification of experience to be a spectral notion, with veridical perception and hallucination at either end, and illusion occupying the middle space. Since illusion involves this mixing and matching of feature veridicality, it is permissible – in cases such as only a single particular feature appearing non-veridically – that the mistake will be so minor that the whole experience is still deemed veridical, however, with a similarly scalar divide obtaining between illusion and hallucination.²²⁷

I have focused predominantly on Fish's account of illusion, but that is not to suggest that it is the only option available to the relationalist. Since relationalism holds that the phenomenal character of perceptual experience is determined *in part* by a perceptual relation of acquaintance, this leaves open to the relationalist the possibility that illusory aspects of a given conscious experience are determined by *non*-relational aspects, which determine *in part* the overall phenomenal character of experience.²²⁸ This observation operates at a higher level than Fish's account, and whilst it is compatible with his view, it neither entails it, nor is incompatible with being satisfied in other ways.

²²⁵ Ibid.

²²⁶ Ibid.

²²⁷ Ibid. (178 – 179)

²²⁸ Soteriou (2016: 188)

At its core, relationalism is a positive view on the fundamental account to be given of veridical experience. Disjunctivism, as has been hitherto discussed, is a negative view brought about to support this in light of the argument from hallucination. By adopting a V vs IH (as opposed to VI vs H) conception of experience, the disjunctivist effectively deals with the argument from illusion too, which may account for the content of Logue's observation that this is what several disjunctivists have done without supporting argument. However, the supporting argument *is* necessary – if we are to treat illusions as kindred to hallucinations, there should be an account of illusion that tells us why we should do this. In the absence of this, it remains an open question how the relationalist should account for illusion.

Since relationalism is a view regarding the *fundamental* account of perceptual experience, it is not incompatible with the occurrence of representation, as the representationalist uses the term. What it is generally taken to be incompatible with is that representation does the explanatory heavy lifting. If it is true that relationalism is a view about the fundamental account of *veridical* perception, however, then it is not clear that representation does not play some role in accounting for the illusory aspects of experience. Adopting Fish's threefold distinction of physical, cognitive and optical illusions, the relationalist could even hold that representation accounts for the non-veridical aspects of illusion, whilst their preferred articulation of relationalism accounts for its veridical elements. This does not seem to violate the tenet that the acquaintance relation is what fundamentally accounts for veridical experience. I do not think this altogether implausible, for what it is worth, but concede that it may be uncomfortable for hard-line austere relationalists to accept and will not argue for it further here.

A final observation on how the relationalist might account for illusion comes from the claim that the best wholly relationalist explanation of illusion may not involve what one is perceptually related to, but rather *how* one is related to it. This is the position occupied by Brewer's account of illusion, discussed above in §2.2. According to this VI vs H position, when it illusory appears to

a subject *S* that an object *o* is *F*, when it is in fact not *F* but *G*, what has happened is that *o* bears visually relevant similarities to paradigm exemplars of *F* to *S* *from the standpoint* that *S* is perceptually related to *o*. *S* registers this and as a result forms the mistaken belief that *o* is *F*. It is worth noting that Brewer's view also features the notion of *degraded acquaintance*, whereby a smaller quantity of an object's determinate features form part of the acquaintance relation in question, which can also lead to erroneous judgements as to *o*'s overall nature.²²⁹ Furthermore, there is nothing to prevent accounting for illusion of object *properties* further by appealing to certain hallucinatory elements to experience being 'superimposed' on a mind-independent object that one correctly registers the presence of.²³⁰

There are, if anything, more options available to the disjunctivist regarding illusion than there are regarding hallucination, and this section has been structured so as to provide an outline of the various ways in which both might be accounted for by relationalists who adopt the view. In the following section, I turn to discuss a different, less metaphysical variety of disjunctivism, which the relationalist might nevertheless appeal to in order to uphold their view in the face of the problems posed to their view by non-veridical experiences.

2.4.2: Epistemological Disjunctivism

An assumption underpinning the metaphysical brand of disjunctivism put forward by Martin and endorsed by others, is that the difference between veridical perception and hallucination is to be located at the level of psychological states.²³¹ One way to potentially avoid the objections to that

²²⁹ See Brewer (2017: 224 – 225)

²³⁰ Soteriou (2016: 188); see Macpherson & Batty (2016) for support of the view that a subject's experience can be misleading with regards to object and properties, and that there is a double dissociation between these sorts of error. See Brewer (2011, 2017) for more on the notion of 'superimposed' hallucination being used here.

²³¹ Brewer (2011) also accepts a view much like this.

view, then, is to argue that the difference between the two sorts of experience is located somewhere else, and then claim that it is *this* difference that allows explaining the two sorts of experiences differently. One such account is put forward by John McDowell, in his (1982), and is commonly known as *epistemological disjunctivism*.

This view advances an alternative, epistemological means of rejecting the Highest Common Factor thesis:

(HCF): When you successfully perceive your environment, the warrant that you have for making judgements about your environment is no better than, no stronger than, the warrant that you would have for making such judgements if you were having a subjectively indistinguishable hallucination.²³²

A rejection of (HCF) commits the epistemological disjunctivist to viewing a subject's warrant for making judgements about their environment as being different in veridical and hallucinatory cases. In summary, epistemological disjunctivism supports this epistemological claim by arguing that in hallucinatory but not in veridical experiences, one's experience falls short of the fact:

*In a deceptive case, one's experiential intake must ex hypothesi fall short of the fact itself, in the sense of being consistent with there being no such fact. [. . . The] object of experience in the deceptive cases is a mere appearance. But we are not to accept that in the non-deceptive cases too the object of experience is a mere appearance, and hence something that falls short of the fact itself [. . .] The idea of a fact being disclosed to experience is in itself purely negative: a rejection of the thesis that what is accessible to experience falls short of the fact.*²³³

In cases of veridical perception only, then, the epistemological disjunctivist claims that what is accessible to experience is the fact itself. This is epistemologically appealing, since it is intuitively plausible that what makes it the case that a subject knows *P*, when it is in fact the case that *P*, just is *P* itself experienced by the subject; this intuition can be neatly accommodated by a view according to which “the content of the experience

²³² Soteriou (2016: 118)

²³³ McDowell (1982: 471 – 472)

is the content of the knowledge.”²³⁴ Conversely, hallucinations are by definition such that it is not in fact the case that *P* (or, in the case of veridical hallucinations, it is not *P* that brings about the subject’s experience even though it is in fact the case that *P*), and consequently it cannot be *P* itself being experienced by the subject that makes it the case that they know *P*. The hallucinating subject does not know that *P*. The content of the experience cannot match the content of the knowledge if there is no knowledge.

Examining the epistemic asymmetry between subjects of veridical perceptions and hallucinations further, epistemological disjunctivism further claims that what the subject has access to in the good case alone is *conclusive perceptual warrant* for believing that *P*.²³⁵ This conclusive warrant is *P* itself, and so is unavailable in the bad case. If this is correct, then it follows that the warrant one has experiential access to in the good case for making judgements about one’s environment *is* stronger than the warrant one has to do so in the bad case. This entails the falsity of (HCF), which is the conclusion that the epistemological disjunctivist wished to establish.

The view must take account of the same particulars of hallucination as the metaphysical variety of disjunctivism, however. In addition to claiming that subjects in the good case have access to conclusive epistemic grounds for judgement for perceptual judgement that those in the bad case lack, the epistemological disjunctivist must acknowledge that when a subject is in the bad situation, they would fail to notice the lack of such epistemic grounds.²³⁶ This is to acknowledge the subjective indistinguishability between the good case and the corresponding bad case. In the bad case, one is deceived with respect to the layout of one’s environment, as well as with respect to the kind of experience one is having.²³⁷

²³⁴ Ibid. (174)

²³⁵ Soteriou (2016: 128 – 129)

²³⁶ Ibid.

²³⁷ Ibid. (133)

It is uncontroversial that what we can know is that things in our environment appear to be a certain way. What this amounts to, according to McDowell's view is a *disjunctive fact*: either an objective state of affairs makes itself manifest to us, *or* a situation in which it is *as though* this is happening occurs.²³⁸ The subject is therefore fallible with regards to which of these disjuncts obtains. A possible objection to McDowell's position is that, if one is only aware of the disjunctive fact, then one does not have the access to conclusive perceptual warrant previously described.²³⁹ However, it can be argued that experience therefore has a *fallible* capacity to provide us with conclusive epistemic grounds, and in situations where this fallibility does not kick in, it actually does provide us with conclusive epistemic grounds. Consequently, once defective and non-defective exercises of that fallible capacity are treated as epistemically asymmetrical, the claim that non-defective exercises of the (fallible) capacity provide conclusive grounds for perceptually-based judgements can be upheld.²⁴⁰

Cases where the subject has conclusive grounds for judgements about their environment highlight why epistemological disjunctivism might be appealing to relationalists. Perceptual experience provides the subject with conclusive grounds for judgements, whilst hallucination and illusion cannot, and so this is an epistemological aspect of perceptual experience for which we must provide an account. Since the epistemological disjunctivist claims that this epistemic asymmetry operates between good and bad cases, they hold that what provides the subject with conclusive grounds for making judgement is the presence in their environment of the mind-independent objects about which they are making judgements. This means that the epistemological aspect of perceptual experience to be explained here is to be explained in terms of the relation a subject stands in with some mind-independent object(s). This view matches well with relationalism, which claims that

²³⁸ Ibid. (134)

²³⁹ Ibid. (135)

²⁴⁰ Ibid. (135)

perception is fundamentally a matter of a subject standing in a relation of perceptual acquaintance to mind-independent objects.

As described, epistemological disjunctivism allows the epistemic asymmetry between the good and bad cases to be explained in terms of relations to mind-independent objects. These mind-independent objects, according to the view, serve as conclusive epistemic grounds for subjects making perceptual judgements about their environment. In this respect, the subject has a reason (epistemically speaking) for judging their environment to be a certain way because of the presence of mind-independent objects. This is an important point, and I shall return to it in Chapter 4.

One of the problems with metaphysical disjunctivism was that its minimal account of hallucination, which is negative and epistemic in nature, is underwhelming given the metaphysical difference between the good and bad cases that the view suggests. As should be clear from the preceding discussion, epistemological disjunctivism makes no claim to veridical and hallucinatory experiences being fundamentally different sorts of psychological states. According to epistemological disjunctivism, the difference between veridical perception and hallucination *qua* experience is epistemic in its entirety. Since this is all that the view then needs to explain, accounting for it in purely epistemic terms is not the lacklustre explanation that it seems to be for metaphysical disjunctivism.

We saw above that what the metaphysical disjunctivist ought to say about illusion is up for debate. Although there was not nothing to be said, the commitments of the view raised the immediate question as to which disjunct of psychological states illusion – being non-veridical but also object-involving – is to be aligned with. The metaphysical disjunctivist must propose and support an answer to this question, and an account must then be provided of how they explain this. A further benefit of epistemological disjunctivism is that it side-steps the need to make this decision. This is because the view makes no claim of a divide in psychological states to begin with. Clearly, there is a difference in epistemic state between the subject

undergoing an illusion and a veridical perceiver, and the view can accommodate this in a manner analogous to its treatment of hallucination.

For epistemological disjunctivists, further explanation of illusory experience need not be couched in terms of their disjunctivist position at all, but can instead be couched in terms of positive relationalist apparatus. It can appear to a subject that a blue car viewed at night is actually green because of the yellow tint of the streetlight overhead and the otherwise poor lighting conditions, for instance. More broadly, this is a return to the point made in §2.2, to the effect that illusion involves a subject standing in a relation of perceptual acquaintance to an object, as in veridical perception, yet the standpoint from which they are related to the object is such that they judge it to be a way other than it is. Returning to the previous example of the blue car at night, its subjective indistinguishability from a green car is thereby explained wholly by the three-place relation of the subject, the car itself, and the standpoint from which the subject experiences it.

It is worth re-emphasising that metaphysical disjunctivism is not precluded from adopting this approach, but this entails the further commitment that there is a difference in psychological state between illusory and hallucinatory experiences. The metaphysical disjunctivist may well accommodate this, so this consideration is not a decisive indicator of which variety of disjunctivism the relationalist ought to accept. That said, it is arguable that epistemological disjunctivism need not make any such commitment, and therefore that the view involves fewer complications than its epistemological counterpart.

I conclude this discussion of the varieties of disjunctivism with two further brief points. The first of these pertains to an important commonality between epistemological and metaphysical disjunctivism. As described above, a feature common to both views is that veridical perception and hallucination differ with respect to the epistemic states of subjects undergoing each sort of experience. More precisely, veridically perceiving subjects possess, whilst hallucinating subjects lack, epistemic reasons for

forming judgements about the nature of their environment that are grounded in their experience. One appealing feature of relationalism explored above was that it can provide an account of what makes both knowledge and thought about the external world possible – claiming that experience plays the grounding role that enables both. Each of the varieties of disjunctivism discussed above maintain this grounding role of experience. It follows from this that regardless of which of the two accounts of non-veridical experience the relationalist adopts, they are still able to uphold that experience plays this role. Since what the veridically perceiving subject has experience of is, on each view, some mind-independent object, it follows that relationalists can maintain that mind-independent objects themselves constitute epistemic reason for a subject to form judgements about their environment iff that subject veridically perceives the mind-independent objects about which they are forming judgements.

Finally, the objector to disjunctivism may attempt to argue against any apparent benefits of adopting epistemological disjunctivism by arguing that the view is itself committed to the metaphysical variety of disjunctivism, and therefore susceptible to the same objections as that view is. In the first instance, the onus is upon the objector to argue this point. Assuming epistemological disjunctivism temporarily, veridical perception and hallucination both involve the exercise of an epistemic capacity, correctly in the former and defectively in the latter. It does not obviously follow from this alone that a subject is in a different psychological state in each of these cases; in other words, it is not clear that the radically different disjuncts of epistemological disjunctivism cannot share a common metaphysical essence of the sort precluded by metaphysical disjunctivism.²⁴¹

According to the epistemological variety of disjunctivism, what reason we have to form judgements on the basis of veridical experience does not fall short of the facts (i.e. of mind-independent objects in our environment), and so veridical experience *directly* – that is: non-inferentially – justifies our

²⁴¹ Pritchard (2012: 24) and Soteriou (2016: 146 – 147) both make this point.

perceptual beliefs. Hallucination cannot do this, and so epistemically the two are importantly different. There is no logical connection, however, between this epistemic directness and the directness that centrally concerns metaphysical disjunctivism. The claim that veridical experience can directly justify our perceptual beliefs is compatible with *any* metaphysical view of perception according to which, after the veridical experience, there are no further evidential requirements imposed on the resultant perceptual belief.²⁴² In this regard the view is compatible with metaphysical disjunctivism, but is also compatible with sense-datum theory, provided that any inference from sense-data to mind-independent objects meets the aforementioned criterion, perhaps by being an unconscious or sub-doxastic inference.²⁴³

On balance, it could nevertheless be argued that it is hard to *motivate* epistemological disjunctivism without appeal to its metaphysical counterpart.²⁴⁴ The objector could then claim that unless the epistemological disjunctivist motivates their position without appeal to metaphysical disjunctivism, their argument goes through. The issue with this approach is that it is not clear that the metaphysical disjunctivist can articulate their position without explaining matters in epistemic terms.²⁴⁵ Their good and bad cases differ *qua* experience just because the former puts the subject in a position to have perceptual knowledge, whereas the latter (according to Martin, at least) does not, and can only be described as a case that is subjectively indistinguishable from this. As we have seen, Martin develops this view and deals with the prospect of causally matching hallucinations in wholly epistemic terms.

If this is correct, then the best our objector could hope for here is a strangely circular outcome for disjunctivism, in which they are right that epistemological disjunctivists cannot motivate their view without appeal to the metaphysical variety of disjunctivism, yet the converse is also true. Given

²⁴² Drayson (2017: 15 – 16).

²⁴³ Lyons (2016)

²⁴⁴ Pritchard (2012: 24)

²⁴⁵ Soteriou (2016: 149)

the above point regarding the logical independence of the views, this seems like an untenable outcome. Further, since there seem to be some considerations counting against the objector's position, and none obviously counting in its favour, I support the conclusion that the success of epistemological disjunctivism should not be linked to that of metaphysical disjunctivism, until such time as further argument against this conclusion is given.

3: *Explanatory Pluralism*

Thus far, I have argued that contrary to what until recently was apparently the orthodox mentality in the philosophy of perception surrounding the relational and representational views, an argument raising concerns or casting doubt upon one does not necessarily support the other. It is not, in other words, a ‘one or the other’ situation. This is because in many cases, an argument advertised as an argument against relationalism or representationalism is, in fact, an argument against the *austere* version of relationalism or representationalism.²⁴⁶ For this reason, it is in principle possible to split the explanatory labour between the relational and representational views and, I submit, also to utilise the arguments which allegedly raise problems for one view or the other in order to set some kind of guiding limitation on exactly how the explanatory labour might be split up.²⁴⁷ Given the distinction between austere and non-austere versions of relationalism and representationalism, the successful articulation of this sort of view would incorporate non-austere versions of each theory. Since neither of these versions commit their proponents to rejecting the alternative view, there is no tension following from the simple tenets of the views being accepted in the construction of the overall theory of perception. Also, since the view to be advanced incorporates (non-austere) relationalism and representationalism based on what it is about perception that each view can explain effectively that the other apparently cannot, then – to the extent that examples of such instances can be provided – it seems that such a view can constitute an explanatory pluralist view that can circumvent the Screening Off Argument.²⁴⁸ Ideally, though, this would be accomplished whilst *also* avoiding the charge of being an ‘anything goes’ relativist position which

²⁴⁶ Where austere relationalism, for instance, is the view that perception is fundamentally relational (in the sense the relational view claims) *together with* the claim that it is *not* fundamentally a matter of representation (see §1.2).

²⁴⁷ The thought here being that the view resulting from following through on Logue’s suggestions would benefit from not being wholly relativist. I return to this point below.

²⁴⁸ Discussed in §0.3.

holds that the preference or exclusion of explanations cannot be achieved on any level. For each individual thing to be explained, such a relativist view would deny that any ranking of potential explanations was actually possible. Taking together the thoughts that (a) the Screening Off Argument might be circumvented by appeal to explanatory pluralism, and (b) such an explanatory pluralist position must not become unrestrictedly relativist, it follows that there must be some means of ranking the suitability of explanations surrounding the issues of perception at hand (commitment of (b)), but also that there cannot be any global exclusion or preference rules operative over all explanations, or perhaps even all explanations of a particular type (commitment of (a)).

The Screening Off Argument, introduced previously, has not yet been sufficiently addressed. Yet it seems that at the very least the reasoning involved in that argument is similar to that which might be used to deny the above contentions that an argument against e.g. austere representationalism is not an argument against the plausibility of a reconciliatory view which has features of representationalism within it. The main task of this chapter will be to commence in the process of dealing with the Screening Off Argument. In summary, the case to be made here will motivate and utilise the position known as explanatory pluralism, together with some other theoretical apparatus and terminology from the literature on explanation, to undermine the claim that finding, say, the representational view a place in our satisfactory philosophical theory of perception renders the relational view explanatorily redundant. Ultimately, this shall be because having the representational view (for example) constitute the *entirety* of our theory of perception is not, as I hope to show, the only plausible way to find a place for the view within our theory.

In what follows, I shall first introduce explanatory pluralism in more detail, as it plays a role of paramount importance in what follows. Next, I observe the areas of discourse in which explanatory pluralism enjoys at least some support in terms of its motivation, and assess the case for its utilisation in resolving the relationalism/representationalism compatibility issue within

the philosophy of perception. Then I shall introduce and discuss some other notions from the literature on explanation, utilising these to further elucidate and augment the broad approach sketched out by Logue and some others. Overall, the chapter begins setting the scene for showing in Chapter 6 that the Screening Off Argument can be overcome.

3.1: *Introducing explanatory pluralism*

The first task ahead is to define and provide some description of the workings of explanatory pluralism. As a first-pass definition of the view, following Bouwel and Weber, I take it that explanatory pluralists – broadly construed – are at least committed to the following two claims:

- 1) There are no general *exclusion rules* with respect to explanations [...]; it is, for instance, impossible to rule out intentional explanation or functional explanations.²⁴⁹
- 2) There are no general *preference rules* with respect to explanations [...]; it is, for instance, unwarranted to claim that intentional explanations are always better than macro-explanations.²⁵⁰

Unpacking what is going on in (1) and (2), immediately it can be seen that these are both negative claims and, more particularly, negative claims which each constitute a rejection of some principled, general rules regarding explanation. In the first case, the rules being rejected are exclusion rules – rules by which a given kind of explanation might be ruled out as an illegitimate explanation *in virtue of the type of explanation it is*. In the second case, the sort of rules being rejected are preference rules – rules by which two or more types of explanations

²⁴⁹ For clarity: the types of explanation here correspond to types of explanations identified in social science, as opposed to whatever they may be taken to mean in the context of academic philosophy.

²⁵⁰ (2008: 168); ellipses in each claim exclude the words “in history and social science” – utilisation of explanatory pluralist methodology in particular sub-disciplines will be focused upon in the next section, so I save discussion of such points for then.

might be arranged in a general preferential hierarchy of explanations, again in virtue of the type of explanations they are.

To clarify, neither (1) nor (2) commit to the respective claims that either it is never possible to discount an explanation as illegitimate, or that it is never possible to prefer one sort of explanation to another. Saying anything that further elucidates how this is possible requires an examination what we might call the more positive commitments of explanatory pluralism. However, to achieve this it will be necessary to introduce and examine some further terminology and concepts from the literature on explanation.

In the broadest sense, what is required first and foremost is the notion of *explanatory context*. To see the centrality of this notion, consider first the words of Schweder, who writes that “the reification of explanation, the belief that some propositions are, once and for all, explanations, is quite mistaken. Only particular statements, utterances or other acts of communication can be explanations, and whether they are so, or not, is contextually determined.”²⁵¹ This, admittedly, does leave what an explanatory context actually is a little up in the air, and I will attempt to shed light on this in what is immediately to follow. On the view to be advanced here, though, this notion of explanatory context is crucial to explanations serving their purpose.

In attempting to elucidate further the notion of explanatory context, Schweder provides a cursory analysis of the nature of explanation itself, and what this very notion implies or presupposes. Her analysis focuses on the format of explanation as suggesting an explanation-seeking agent and an explanation-giving agent or, simplifying slightly, as suggesting a model akin to an explanation-seeking question (henceforth an *es-question*) and an explanation-giving answer (*eg-answer*).²⁵² More broadly still, an *es-question* may be thought of as involving “an

²⁵¹ (1999: 116)

²⁵² *Ibid.* (116 – 117)

inquirer” whilst an eg-answer may be thought of as containing “an *explainer*”, a distinction which appears to map quite neatly onto the possibly more familiar *explanandum/explanans* distinction.²⁵³ By way of abstract example of how es-questions and eg-answers work in practice, we have the following schema:

Es-question: “Why B?”

Eg-answer: “B, because of A” or, equivalently (in the case of causal explanations) “A explains B.”²⁵⁴

Here, B is going to play the role of *explanandum*, and is to be taken as a variable to which is assignable, in the simplest instances, the many ways in which an event may be described. For example, the event in which Saul goes to the shop may be described as (i) Saul walking down a sequence of streets (with emphasis on the route Saul has taken), (ii) Saul going to the shop (as opposed to elsewhere), (iii) Saul going to buy beer, (iv) Saul spending a part of the day doing a certain activity (as opposed to any other activities in which Saul may have potentially engaged), or (iv) Saul going to work, and presumably in several other distinct ways besides. Here though, it seems as though plugging in any of (i) – (iv) into the es-question part of the above schema admits of various sorts of question in turn. Say we plug in (iv); we have the question ‘why did Saul go to the shop?’, with the emphasis on going to the shop as being the particular activity in which Saul engaged. Yet here it seems again that this question, as phrased above, admits of various interpretations that alter based on how the question is asked:

Non-contrastive: Why (as in: for what reason) did Saul go the shop?

Property contrastive: Why did Saul go to the shop, as opposed to going for a walk on the beach?

Subject contrastive: Why did Saul, as opposed to Seth, go to the shop?

²⁵³ Ibid.

²⁵⁴ Ibid. (117)

Time contrastive: Why did Saul go to the shop at time *t*, as opposed to at some other time?²⁵⁵

In their fully articulated forms, the distinctions between these questions are clear, yet they are all perfectly reasonable interpretations of the original ‘why did Saul go to the shop?’ question. The latter then, it seems, can be asked as an es-question that seeks different explanations in different instances and – in this way – what counts as a *relevant* eg-answer to it can vary. In other words, different facts are relevant to different interpretations of the same es-question, and different inquiring agents may be seeking different answers despite uttering identical words in asking their es-questions.

As Schweder notes, the notion of *reference class* is a useful one to bring into play here, and can help elucidate the point: supposing the inquirer is asking a subject contrastive es-question – and wishes to know why Saul went to the shop, as opposed to anyone else – it seems that a fitting reference class would be all the people who live in a particular household, namely Saul’s household.²⁵⁶ Alternatively, suppose that a different inquirer uttering the identical words ‘why did Saul go to the shop?’ is in fact asking a property contrastive (p-contrastive) question, wishing to know why Saul went to the shop, rather than going running or going to the beach. Here the relevant reference class is plausibly activities: going to the shop in question, as well as the various alternative activities in which Saul may have engaged.

With the notion of reference class in place, it is more readily apparent that the same eg-answer may be satisfying to one inquirer but not to another. Taking the above pair of es-questions, suppose the agent who asked the subject contrastive es-question is offered the latter p-contrastive eg-answer. The relevant reference class for the es-question being answered was a certain set of people (or perhaps the reasons that guide their actions), yet the es-question to which the p-contrastive eg-answer being given is associated has as its reference class activities. It is hopefully uncontroversial that these two

²⁵⁵ Here I am adapting an analysis of sorts of es-questions from Bouwel and Weber (2008: 170)

²⁵⁶ Schweder (1999: 118); or, perhaps more precisely, the reasons governing the actions of that set of people.

reference classes are mutually exclusive: no activity is itself a person; no set of people is a set of activities. We see, then, that certain eg-answers are simply not suitable as answers to certain es-questions.

We also see, contends Schweder, that not only does the particular es-question being asked provide a clue as to where to look for an answer (“it seems that *the answer*, to some extent, is *anticipated* in the question”), but also provides some insight as to the epistemic state of the inquirer – their background beliefs and suppositions – such that certain (possibly more coarse-grained or explanatorily prior) aspects of an eg-answer can be identified as being not in need of explanation for the agent asking the es-question.²⁵⁷ In general, the thought here is that “the es-question encodes several important contextual features. “It comprises *both* the inquirer’s expectations as to what a relevant eg-answer should be, and his background knowledge and presuppositions.”²⁵⁸ On the whole, then, explanation is to be seen as being comprised by both question *and* answer, and “the explanatory value of an eg-answer can only be evaluated relative [to] an es-question.”²⁵⁹ Alternatively, this point may be expressed as stipulating that “[a]n answer to an explanation-seeking question will be adequate for an explainee only if the *explanans* has the expected relevance relation to the *explanandum*.”²⁶⁰

In addition to these important features of es-questions, eg-answers can also be treated to further analysis. For the sake of simplicity, the following few paragraphs conduct this analysis whilst focusing on causal explanations, as I take these to render what follows most readily understandable. Of course, there are other sorts of legitimate explanations besides causal ones, and in providing causal explanations in the following analysis, I do not mean to suggest that *only* casual explanations can constitute an eg-answer.

Additionally, I take it that something like what is discussed below can be equally applied to other types of explanation besides causal, as these too can

²⁵⁷ Ibid (118)

²⁵⁸ Ibid. (119; emphasis original)

²⁵⁹ Ibid.

²⁶⁰ Bouwel & Weber (2008: 171)

be thought of as involving es-questions, eg-answers and explanatory contexts.

As identified in the schema above, an eg-answer will generally be of the form “B explains A”. Contained therein though, contends Schweder, is a counterfactual claim: the eg-counterfactual. Essentially, the speaker of the utterance “B explains A” is purportedly saying that, had B not occurred, A would not have occurred either.²⁶¹ This claim is supported by considering how we would go about disputing that B explains A: if it can be shown that A would have occurred irrespective of whether B occurred, then this constitutes a reason to doubt the claim that B explains A. That said, eg-counterfactuals work in tandem with the notion of a reference class, discussed above. This serves to rule out various irrelevant (albeit true) considerations from a proposed explanation; for instance, it is true that being born is a necessary condition on one’s dying – if one had (counterfactually) not been born, then one could not die – yet in almost no instance of the es-question “why did they die?” would an appeal to their being born be considered a satisfactory eg-answer.²⁶²

As well as ruling out explanatorily irrelevant factors, the reference class also holds constant various factors such that the *explanatory* factor (the one which is most relevant/active in the explanation) can be more readily identified. This explanatory factor, in turn, is the one that is taken to be most likely to render the eg-counterfactual true, given background presuppositions and prior knowledge (i.e. given one’s epistemic state).²⁶³ It is this greater likelihood of making the eg-counterfactual true that, Schweder contends, makes a subject more likely to select that particular explanation, rather than any of the viable alternatives, as the eg-answer to the es question. It is suggested that according to the pluralist position:

In general, we should select the content of our explanation in such a way that it is adequate relative to our motivation for asking the question (that is, relative to our

²⁶¹ Schweder (1999: 119)

²⁶² Ibid. (121)

²⁶³ Ibid. (122)

epistemic interests). This strategy sometimes results in explanations containing remote causes, sometimes in explanations containing proximate causes. So we have a pragmatic reason for being pluralist in the remote causes [e.g. being alive as a cause, albeit in many cases an extremely remote one, of death] vs. proximate causes [say, liver cirrhosis as a cause of death] debate: depending on the context, one or the other type of explanation is the best relative to our motivation for asking the question.²⁶⁴

In the above quotation, the proximate cause is going to be much more desirable as an explanation than the remote one. The thought that proximate causes *always* function as better explanations, though, is precisely the sort of preference rule that explanatory pluralists deny can effectively govern our preference with respect to explanations. By way of a concrete example, consider the memory-degenerative complex of conditions known as Wernicke-Korsakoff complex: comprising a co-presentation of Wernicke's encephalopathy and Korsakoff syndrome, it is believed that the reason for their reliable co-presentation is that each is caused by thiamine deficiency, which in turn (in the 'Western world') is most reliably caused by alcoholism.²⁶⁵ Here, then, thiamine deficiency serves as a proximate cause, whereas alcoholism – being prior to and longer-running than thiamine deficiency – serves as the remote cause. The explanatory pluralist will conclude that the preferable explanation here will depend upon one's interest; with examples of potential interests being (i) reliable intervention (i.e. providing treatment to those presently suffering from Wernicke-Korsakoff complex), (ii) preventing the development of the complex in the later lives of people, or possibly (iii) to understand the difference between two populations to the effect that one of them seems more reliably to be affected by the complex whilst the other does not.

In the case of reliable intervention, manipulating the proximate cause (i.e. administering thiamine) is more effective than attempting to treat alcoholism, as there are other causes of thiamine deficiency (and, therefore of Wernicke-Korsakoff complex) besides alcoholism, such as eating

²⁶⁴ Bouwel & Weber (2008: 173)

²⁶⁵ Gervais (2014: 7)

disorders.²⁶⁶ Conversely, with respect to the case of prevention, treating the alcoholism is likely the better choice, as prediction of those at risk of contracting Wernicke-Korsakoff complex can be accomplished in a more timely fashion by focusing on patterns of behaviour that further precede the undesirable outcome – namely alcohol abuse.²⁶⁷ Finally, with respect to explaining the differences between the two populations, citing a difference in their respective average thiamine levels will do little to demystify the *explanandum*, whereas there is a chance that examining their respective propensities for alcohol abuse has the potential to be much more informative.²⁶⁸

We see, then, that according to explanatory pluralism, the explanation to be preferred depends on our interests. This comes at the cost of denying that there are any privileged levels or types of explanation, at least that aren't relative to our epistemic states and interests. However, this is a far cry from saying that anything goes on such an account – there can nevertheless be guidelines, or procedural heuristics that can direct us towards the likely preferable explanation *relative to our interests and motivations in asking the es-question*. For instance, taking cases (i) – (iii) in the preceding paragraph, we might, for example, devise the following guidelines:

- A) If your interest is reliable intervention, focus on the explanation citing the proximate cause.
- B) If your interest is prevention, focus on an explanation citing a cause distant enough to allow for timely predictions.
- C) If your interest is to understand a contrast between two states of affairs, focus on an explanation citing a cause distant enough to be less mysterious than the *explanandum*.²⁶⁹

Guidelines such as these are all conditional statements, each beginning with the antecedent “if your interest is...”, rendering them sufficiently contextualist such that they are compatible with the tenets of explanatory

²⁶⁶ Ibid.

²⁶⁷ Ibid.

²⁶⁸ Ibid.

²⁶⁹ Ibid.

pluralism, yet also provide some restrictions and (conditionalised) guiding principles, such that they allow the explanatory pluralist to avoid any unwanted charges of unrestricted relativism with respect to explanation.

3.2: *Contextualising explanatory pluralism*

By way of providing some context of how my sketch of explanatory pluralism fits into the wider literature on explanation, credit should perhaps first be given to van Fraassen's *pragmatic theory of explanation*. It is here that we find first the explicit distinction drawn between explanations and propositions or statements, in the sense that the former is *essentially* the answer to a why-question, and therefore a theory of explanations constitutes a theory of answers to why-questions.²⁷⁰ According to van Fraassen's theory of explanation, if one takes a question – for example 'why is this conductor warped?' – the questioner has implied that the conductor is warped (call this proposition the *topic* of the question) and is seeking a reason.²⁷¹ The question will also have a *contrast-class*, or a set of alternatives to the topic (e.g. that *this* conductor warped and not *that* one, or why this one warped instead of retaining its shape – distinct, more specified legitimate ways of interpreting the original question). Finally, we have a relation of explanatory relevance: pertaining directly to the respect in which the request for information was sought in the form of the initial question, this is highly context sensitive and will determine what does or does not count as a possible explanation.²⁷² Sticking with the warped conductor, the respect in which a reason for the warping is requested might pertain to events that rendered the warping possible, which allows as relevant human error, unwanted moisture being present, and witchcraft (as the possibly relevant answers are evaluated *after* they are all gathered).²⁷³ Alternatively, the various conditions leading up to the warping may be understood and in fact

²⁷⁰ van Fraassen (1988: 137 – 138)

²⁷¹ *Ibid.* (143)

²⁷² *Ibid.*

²⁷³ *Ibid.*

what is sought is information as to the immediate catalyst to the process (say the presence of a magnetic field).²⁷⁴

Summing up this latter idea of explanatory relevance, van Fraassen writes:

*In a given context, several questions agreeing in topic but differing in contrast class, or conversely, may conceivably differ further in what counts as explanatorily relevant. Hence we cannot properly ask what is relevant to this topic, or what is relevant to this contrast-class. Instead we must say of a given proposition that it is or is not relevant (in this context) to the topic with respect to that contrast class.*²⁷⁵

Here the connection between his view and the explanatory pluralist one sketched above becomes clear: the conception of explanation as a question-answer model, which admits of various legitimate answers whose very legitimacy is context sensitive, relative to the purposes of the initial request for information, is something that the two views have very much in common. Rather than viewing explanation as a two-place relation obtaining between theory and fact, they each view it as a *three*-place relation obtaining between theory, fact and *context*.²⁷⁶ The role of the latter cannot be understated, as what exactly is being requested by means of the question ‘why is it the case that P?’ varies with context, as does the background theory and data relative to which the question is evaluated, *as does* what part of this background information is used to evaluate how good an answer is, *qua* being an answer to the initial question!²⁷⁷ Saying that a theory can be used to explain a fact, then, is invariably elliptic for something like: “there is a proposition which is a telling answer, relative to this theory, to the request for information about certain facts (those counted as relevant for *this* question) that bear on a

²⁷⁴ Ibid.

²⁷⁵ Ibid.

²⁷⁶ Ibid. (153)

²⁷⁷ Ibid.

comparison between this fact which is the case, and certain (contextually specified) alternatives which are not the case.”²⁷⁸

Ultimately, the point here is that according to both views there is no explanatory power as such: there is simply explanatory power relative to certain contexts, in turn determined by the very request for information involved in the initial question. Comparatively, this illuminates how explanatory pluralism can surmount the Screening Off Argument, for it is not too large a stretch to conclude that if – according to one’s theory of how explanations work – there is no explanatory power *simpliciter*, then there equally may not be explanatory redundancy *simpliciter*, but rather only explanatory redundancy when two explanations are in competition *relative to the same context of explanatory relevance*, and one then clearly overshadows the other.

For all their similarity, van Fraassen’s view does differ from that sketched above insofar as it proceeds through some fairly formal means to express how relevant explanatory factors are to be determined with respect to the context of the request for information, and to express how to evaluate these factors. I do not intend here to address the merits or drawbacks of his approach, opting instead to leave open precisely how best these tasks should be done. The purpose of this discussion, rather, was to set in literary context the sort of explanatory pluralist view I have sketched out. There are other views in the literature with which the above sketch (which I broadly attribute to Schweder) coheres well.²⁷⁹

This is not to suggest that the methodology I have laid out is the only articulation of explanatory pluralism. Indeed, Mitchell *et al.* draw distinctions between different sorts of pluralism, such that it is in fact quite clear what position the view I am suggesting adopts. They draw two distinctions: one between *competitive* pluralism and *compatible* pluralism;

²⁷⁸ Ibid.

²⁷⁹ See Gervais (2014) for a contemporary outlining of an explanatory pluralism extremely similar to the one I am using. Kendler (2005) and Mitchell *et al.* (1997) also provide illuminating discussion on alternative pluralistic accounts.

the other between *isolationist* pluralism and *integrative* pluralism. To the first distinction: the first of these – competitive pluralism – is taken as a view that “any fallibilist with respect to scientific truth must accept”, as it is simply the view that a pluralism of competing hypotheses being entertained makes for more rapid progression in terms of spotting flaws in theories and getting closer to correct ones.²⁸⁰ Whilst I do accept this point, it is not merely in this respect that the view I am endorsing qualifies as pluralist. Compatible pluralism, conversely, is far closer to the view I am endorsing. This flavour of pluralism appeals to “distinct levels of analysis”, in turn allowing for multiple distinct hypotheses to be accepted compatibly on different levels of analysis.²⁸¹ Consequently, “different *questions* invoke different explanatory schemata,” and “answers to questions at the different levels represent compatible components of a pluralistic, multidimensional body of knowledge.”²⁸² In other words, as Sherman puts it, “competition between alternatives appropriately occurs within and not among levels.”²⁸³

Addressing the second distinction, isolationist pluralism does also employ a ‘levels of analysis’ framework, yet limits the interactions between the theories that are being appealed to at various levels of a given explanatory domain.²⁸⁴ This isolation of theories that minimises their interaction may be good for modelling certain explanations without fear of unexpected variables impacting results, but again it is not the sort of pluralism to which I wish to subscribe. Integrative pluralism, on the other hand, has no such restriction, and is based on the observation that, when one moves away from idealisations and into actual application:

[O]ne can immediately see that causal models that provide answers at different levels are indeed related. Thus, although pluralism is to be defended, it is not the pluralism of questions and the consequent independence of answers, but rather a

²⁸⁰ Mitchell *et al.* (1997: 117)

²⁸¹ *Ibid.*

²⁸² *Ibid.* (118)

²⁸³ (1988: 616)

²⁸⁴ Mitchell *et al.* (1997: 121)

*pluralism of models of causal processes that may describe contributing factors in a given explanatory situation.*²⁸⁵

Thus, on these distinctions, the view I wish to endorse might be referred to as *compatible, integrative explanatory pluralism*: a view which, though admitting of partial autonomies of theories, concepts, and particularly levels of explanation, ultimately advocates a focus on how these distinct parts of a whole fit together.²⁸⁶

With this summary of how explanatory pluralism works now laid out, I can turn in the next section to attempting to motivate the notion that such a view on explanation can or should be applied to the debate in the philosophy of perception that is the central question of this thesis: namely the matter of compatibility between relationalism and representationalism.

3.3: *Motivating explanatory pluralism*

So far in the above discussion, none of the resources which have been appealed to in order to elucidate explanatory pluralism have, in fact, constituted efforts to argue for the prudence of adopting an explanatory pluralist approach within the philosophy of perception.²⁸⁷ In lieu of any such efforts, an attempt to motivate the utilisation of explanatory pluralism in this way – which shall be required in some form or another for my effort to overcome Screening Off Argument – will need to take a form distinct from simply appealing to work already done to that effect. The way I see it, this leaves a live option open which is composed of two stages. In the first of these stages, one appeals to efforts made to motivate explanatory pluralism in other (sub-)disciplines, with a view to identifying just *what* it is that allows

²⁸⁵ Ibid. (121 – 122)

²⁸⁶ Ibid. (149 – 150)

²⁸⁷ Bouwel & Weber (2008), suggests adopting explanatory pluralism in history and the social sciences, for instance. Schweder (1999), conversely, is a more general analysis of causal explanations. Other sources to be discussed in what follows also do not constitute examples of attempts to motivate utilising explanatory pluralism in the philosophy of perception, though Dale *et al.* (2009) discuss using such an approach in cognitive science, Vreese *et al.* (2010) in medical sciences, and Kostić (2017) discusses using a comparable view in the study of consciousness (albeit with respect to the so-called ‘hard problem’ of consciousness).

the claim that explanatory pluralism is motivated in those areas to pass muster. In the second stage, one argues that the use of explanatory pluralism within the philosophy of perception can be similarly motivated, and that no defeaters emerge from the differences in (sub-)discipline involved.

3.3.1: Explanatory pluralism in history and the social sciences

Turning again to the work of Bouwel and Weber, cited above, we see that the examples they utilise are to the effect of motivating the use of explanatory pluralism in history and the social sciences. Of the various accounts to be discussed in this set of subsections, this is perhaps the one at greatest remove from the philosophy of perception, hence I begin here.

Consider cholera. More specifically, two hypothetical yet realistic patterns of outbreaks thereof, such that both Koch City and Miasma City each reliably experience outbreaks every decade or so, after a summer of intense rain; one year (call it year X), it is stipulated that Koch City denizens remain healthy after a summer of intense rain whereas Miasma City's population experiences a cholera outbreak once more.²⁸⁸ Supposing we wish to explain this change in outbreak pattern by way of a macro-level causal explanation, as are often relevant to history and social sciences (where one may wish to explain one social phenomenon by appealing to a distinct social phenomenon that caused it), we can once again draw a distinction between plausible causal explanations along the remote/proximate cause distinction.²⁸⁹ Keeping in mind this distinction, two possible explanations can be entertained:

- 1) (Proximate cause): The population of Koch City remained healthy, while Miasma City was hit by cholera, because *cholera bacilli* were produced on a

²⁸⁸ Bouwel and Weber (2008: 172)

²⁸⁹ Ibid. (171 – 172)

large scale in Miasma City, while their number remained very limited in Koch City.

- 2) (Remote cause): Koch City built a new sewerage system after the previous outbreak in the year X-10. Miasma City does not have a new sewerage system.²⁹⁰

Supposing that each of these are indeed viable macro-level, causal explanations of the difference in cholera outbreak patterns in the year X, it would seem that we would ideally have some way of selecting one as preferred. Yet by hypothesis each of them satisfies the initial constraint on type of explanation. One option for how to proceed here would involve principled rules for selecting a preferred explanation based simply on whether the cause cited is proximate or remote. Explanatory pluralists, however, will not wish to buy into this strategy, as it goes against the very core of the explanatory pluralist position. Another way to go, which would avoid this outcome, would be to base the choice of preferential explanation on the inquirer's interests in attaining the explanation. One question that can be asked, in this vain, is *why* we wish to explain the difference in the pattern of cholera outbreaks. Here, a very plausible answer would be 'to help the people of Miasma City' (hence the explanation would serve a therapeutic function).²⁹¹

If the foregoing analysis of the situation is correct, we can refer to our heuristic, interest-relative guidelines introduced in the preceding section when the case of Wernicke-Korsakoff complex was being discussed. The cholera example seems a case such that what is to be explained is a difference between two states of affairs - namely: the state of affairs obtaining in the year X-10 and that obtaining in the year X. Referring to our guidelines, what these suggest we should be focusing on is a cause remote *enough* such that it is less mysterious than the *explanandum*. Notice, though, that the emphasis on "enough" in the relevant guideline suggests

²⁹⁰ Ibid.

²⁹¹ Ibid. (172)

that it is in principle possible that the proximate cause in this instance is sufficiently remote for the purposes of satisfying this demystifying criterion. What must be done from here, then, is a comparison of the two explanations to work out which *best* suits the therapeutic interests of the inquiring agent.

Considering first the proximate cause explanation, it seems that this leaves open the question of *why* there was a reduction in the production of *cholera bacilli* in Koch City but not in Miasma City. Conversely, the explanation citing the more remote cause of Koch City's new sewerage system is such that it cites a human intervention present in one case and absent in the other and that is precisely what, in Bouwel and Weber's view, allows this explanation of the two to serve the therapeutic function relevant to the interests of the inquirer.²⁹²

In this case, it seems that what motivates explanatory pluralism is simply the thought that a means of selecting the preferred of a plurality of viable explanations was required, for there was more than one explanation satisfying the macro-level causal criterion, and there was nothing about the citation of more or less proximate causes in this plurality of explanations that intrinsically offered a decision-making procedure in this regard. Rather than this being the correct account of what motivates explanatory pluralism in this case, however, perhaps there is something about history and social sciences that are the real motivating force. Though it is my contention that this is not the case, this claim will be more plausible once we have some basis for comparison, and so in the following subsection I turn to examine motivations for embracing explanatory pluralism in areas distinct from history and social science: namely, medical science and cognitive science. Thereafter, the hope is that this will constitute a broad and varied enough

²⁹² Ibid.; Bouwel and Weber go so far as to contend that "in every context where the proximate causes do not relate to human interventions, only explanations in terms of more remote causes can have a therapeutic function." This is a strong and interesting claim, however not one I shall defend here.

set of subject areas, such that claims about motivating explanatory pluralism in the philosophy of perception do not seem quite so *ad hoc*.²⁹³

3.3.2: Explanatory pluralism in psychiatry

One area of study in which explanatory pluralism is arguably motivated that connects centrally to the mind is that of psychiatry. I take this to be a useful area to examine since although related to the mind it is at some remove from the central focus of this thesis – our understanding of the perceptual aspects of our conscious experience – and is thereby less entangled with potential objections and controversies that arise in the latter, which shall be tackled later. Whilst my overall focus in this thesis is that of philosophical theories adopted to further our understanding of the nature of our perceptual experience, psychiatry instead focuses on the aetiology of mental disorders and diseases: it is directed at uncovering and studying the causal explanations of these disorders and diseases. In this respect, it fits quite neatly into the current dialectic of showing that explanatory pluralism can be motivated with respect to causal explanations in varying fields of study. In what follows of this section, I shall be discussing the work of Kendler, who claims that explanatory pluralism is indeed motivated in psychiatry. I shall first offer some general remarks and exposition on the contention that explanatory pluralism is motivated in the field of psychiatry, before discussing several examples highlighted in the literature to illustrate this point.

Considering first, then, the broad notion that explanatory pluralism is appropriate for deploying in the field of psychiatry, it is worth considering that to which such a claim would amount. Recalling earlier discussion on the tenets of explanatory pluralism – particularly the two negative claims central to the view – it should be clear on the most general level that adopting an

²⁹³ Brewer, for instance, argues that reconciliatory theories of perception, as I am understanding them here, are ultimately “unnecessary and unmotivated.” (2017: 221 – 222) I take part of the challenge ahead of me, therefore, is to dispel any concerns about lack of motivation or *ad hoc* acceptance of such views.

explanatory pluralist methodology here would amount to the rejection of general rules for both (i) explanatory preference and (ii) explanatory rejection, in just the same way as we have seen with social sciences. In other words, the explanatory pluralist psychiatrist will reject the notion that any account offered of psychiatric disorders is such that there are any hard, fast rules for selecting the sort of explanation to be preferred *on the basis of the sort of explanation under consideration*. Likewise, on the same basis, there will on this view be no comparable hard, fast rule for rejection of these explanations.

As before, this is a far cry from saying that we have *no* way to sort competing explanations into a hierarchy of preference, it is simply to say that there is no way for this to be accomplished by only looking at the sort of explanation that it is; other factors such as explanatory power, particularly explanatory power based on that which we would like to explain, will be required in order to achieve this successfully. Similarly, the criteria for rejecting an explanation relevant to the field of psychiatry will extend beyond mere examination of the kind of explanation under consideration. As I hope to show, the lie of the land here is no different to social science, or to the area of medical science to be examined below, and selection of our preferred explanation will require careful consideration of the es-questions being asked so that the most appropriate eg-answer can be selected in each case.

If the general tenets of explanatory pluralism involve a rejection of rules of preference and exclusion of explanations, then what this will amount to for explanations of matters of psychiatry in general terms – in the same way we saw it did for social sciences above – is a collection of smaller explanations at various levels, which together combine to most heighten our understanding of the salient *explananda*. The relevance of one or other explanation will vary according to explanatory context which, as we saw earlier, is determined by the es-questions at hand, that is: they are determined by precisely that which we wish to explain.

In support of this broader picture of explanation in psychiatry, here is Kendler:

Our current knowledge, although incomplete, strongly suggests that all major psychiatric disorders are complex and multifactorial. What we can best hope for is lots of small explanations, from a variety of explanatory perspectives, each addressing part of the complex etiological processes leading to disorders. It will be particularly challenging to understand how these many different small explanations all fit together.

In grieving for our loss of big explanations, we similarly have to give up our hope for simple, linear explanatory models. It will not be "A→B→C→D." Etiological pathways will be complex and interacting, more like networks than individual linear pathways.²⁹⁴

The description in the ultimate sentence of this quotation, which describes aetiological pathways as being like networks is particularly telling. Networks, in the broadest terms, are composed of interrelating nodes which are responsive to matters extraneous to themselves. It is an interesting comparison to draw, as it suggests that on this picture of psychiatric explanation, the explanations that may be adopted to account for the causes of various mental problems may themselves vary due to factors extraneous to those very aetiological pathways. Cohering with earlier discussion on explanatory pluralism and its application to social sciences, one such factor may legitimately be the interests of the inquirer: precisely what it is that one wishes to understand about psychiatric conditions, in other words, may be precisely one of the factors that leads to different explanations being adopted.

Whilst it is all well and good to discuss matters in broad terms, without examples of different levels of explanation being incorporated into the study of psychiatry, it will be difficult to persuade anyone of explanatory pluralism's relevance to the field. Consequently, I now examine some examples that illustrate the preceding point that explanatory interests

²⁹⁴ (2005: 435)

dictate the most relevant explanation in psychiatry just as we saw they did with respect to inquiries in the social sciences.

Example 1: Kathy, a young psychiatrist, is asked by a distressed parent to consult with her about her son, Brian, who has decided to leave a career in science to enter the priesthood. The upset parent insists that Kathy order a brain scan to find a way to change his decision. “There must be something the matter with his brain, doctor. How could he throw away such a promising scientific career?” Kathy sees the young man, who appears thoughtful and mature, and he describes the deep satisfaction and inspiration he feels in the Catholic religion. He understands the possible hardships ahead of him but feels he is making the right decision. Kathy tells the parent that she is not going to order a magnetic resonance imaging scan. There is no evidence, she states, that there is anything the matter with his brain, and no interventions that would act directly on his brain are indicated in this situation. She feels that he has reached his decision in a reasonable way, but the mother should feel free, if she wants, to try to argue her son out of his decision.²⁹⁵

In this example, we see demand for one manner of explanation, namely: a magnetic resonance imaging (MRI) scan, being rejected by our psychiatrist Kathy, in favour of another: Brian’s own, doxastic-level reasoning processes. This is important because the two operate at different levels of explanation. Brian’s mother assumes that the correct level of explanation here is that of sub-doxastic processes operating in Brian’s brain unbeknownst to him, hence her insistence on the MRI. However, in Kathy’s expert psychiatric opinion, there is no evidence that the problem (if it is to be called that) can be explained by factors operating at such a level. Tangentially, this touches on the important point that just because we would like or, as inquirers, may suspect that the best explanation occurs on a certain level of explanation, does not mean that the best explanation on offer *does* operate on that level. Based on what Brian’s mother wishes to have an explanation for (i.e. based on her explanatory interests, or the es-question to which she seeks an eg-answer) – namely her son’s sudden change of behaviour and interests – Kathy can have a discussion with Brian and effectively rule out the efficacy

²⁹⁵ Ibid.

of an explanation of his behaviour that operates at the level that the findings of an MRI would provide.

This is, of course, different to saying that an MRI would not be informative: a brain scan could reveal many things that would provide *some* explanation of Brian's behaviours, however these crucially would fail to provide the best eg-answer to his mother's es-question. Since, however, the findings of MRI scans are unquestionably relevant to the explanations adopted by psychiatrists to account for various other psychiatric conditions that they might have to deal with in other patients, it is readily arguable that different sorts of explanation might be preferred. This satisfies the explanatory pluralist tenets that merely looking at the sort of explanations they are is insufficient grounds for either ruling out or preferring one explanation over another.

In this case we see that what allows explanations based on MRI findings to be excluded as the best candidates for providing Brian's mother with an eg-answer simply is the nature of her es-question itself. Whilst she broadly seeks some explanation of his change in mentality, and an MRI would doubtless provide *some* explanation of this, it is concluded that actually a more informative account could be gained by Brian's mother discussing the matter with him further, as this level of explanation is simply more salient than the lower-level mental functioning into which the MRI findings would provide insight (though which *would* be the most salient if her overall goal was to heighten her knowledge of her son's mental processes at the sub-doxastic level). Consequently, it seems, we have here an instance of the tenets of explanatory pluralism being satisfied in the study of psychiatry.

To help illustrate the point further, I offer next an original example featuring two patients with identical physiological symptoms. The patients differ substantially in psychiatric assessment, in that where to look for the most useful explanations of their physiological symptoms varies:

Example 2: Kathy is called to psychiatrically assess two patients who have recently suffered panic attacks. One of the patients, Bill, suffered his panic attack whilst mountain climbing with friends; the other, Nikita, suffered theirs in the middle of a crowded supermarket. Following discussion with each patient, Kathy rules that Bill's panic attack was brought about in response to a near-fatal accident that occurred during the climb, and that the attack was a response to an overload of stress in an otherwise healthy individual. Nikita's panic attack, meanwhile, is deemed by Kathy to be the result of an underlying anxiety condition, which Nikita in turn reported to have flared up due to feeling that many other shoppers in the crowded supermarket were staring at them due to their appearance as an openly trans woman.

Here we see that different sorts of explanation may be utilised in psychiatry even when explaining identical proximate symptoms. In the case of Bill, Kathy appeals merely to factors of his environment – that he was in an extremely high-stress situation immediately leading up to his panic attack. In the case of Nikita, Kathy's assessment does also involve environmental factors – the busy nature of Nikita's environment being deemed a significant part of the explanation of their panic attack – yet also to factors extraneous to merely their environment, in the form of extant mental health issues, themselves explained in terms of social stigma attached to Nikita's gender identity.

In the first instance, the environmental factors are sufficient to serve as where to look for the most salient explanation of his panic attack; yet in Nikita's case, although environmental factors are part of the aetiological pathway that explain the occurrence of their panic attack, finding the most useful explanation of its occurrence requires looking elsewhere. Were Kathy to have adopted a generic preference rule according to which the occurrence of a panic attack (here the *explanandum*) entailed that explanations focused on the patient's environment were the best explanations *based on their focus on that patient's environment*, she would have missed the psychiatric relevance of other factors in determining the cause of Nikita's panic attack.

In order to motivate explanatory pluralism, recall that we need it to be the case that one is required to pick between several explanations, yet lacks a

principled rule that would facilitate this in virtue of the sort of explanations these explanations are. Instead, the preferred explanation is decided based on what which one wishes to explain: whichever explanation provides the most satisfactory eg-answer to one's ultimate es-question is the preferred one. The above example exhibits this well, highlighting that whilst environmental and biological factors may be most useful in explaining Bill's panic attack, far more relevant to ultimately explaining the underlying causes of Nikita's are psychosocial factors. As Kendler puts it:

While certain psychiatric symptoms may be pathological at a basic biological level (e.g., hallucinations), many symptoms are dysfunctional only in certain contexts [. . .] Since many psychiatric disorders include, by definition, some degree of psychosocial dysfunction, explanation at the level of biology alone is unlikely to be sufficient.²⁹⁶

We see here, then, that a pluralist account of explanation in psychiatry can be well motivated. In each of the above examples, acceptance of the view that only one level of explanation is relevant to psychiatric inquiries is simply insufficient to successfully provide the most helpful eg-answers available to the legitimate es-questions that psychiatrists might pose. Focusing on merely biological factors, in the first example, is not going to provide Brian's mother with the answers she seeks, despite her beliefs to the contrary. In the second example, explanations accounting for environment as well as biology are shown to also provide sub-optimal answers when attempting to ascertain the underlying cause of Nikita's panic attack – an es-question that ultimately requires an eg-answer that accommodates the psychosocial aspects of their underlying anxiety disorder.

This concludes my analysis of explanatory pluralism's motivation in psychiatry. In the following section, I shall consider one final area in which explanatory pluralism is arguably motivated – medical science – before

²⁹⁶ Ibid. (437)

progressing on to determine whether this pluralist methodology is similarly well motivated in the philosophy of perception.

3.3.3: Explanatory pluralism and medical science

I turn now to examine the motivation behind utilising explanatory pluralism within medical science. Before delving into their specific examples, though, it is worth noting what Vreese *et al.* say in more abstract terms about using explanatory pluralism in this context:

Suppose we wonder why x has property P at time t. Different, more specific questions— motivated by different epistemic interests—can underlie this general explanation-seeking question, even if all questions are assumed to be requests for causal explanations. For instance, why does x have property P, rather than the more desirable property P'? Is the fact that x has property P the predictable consequence of some other events? Is the fact that x has property P caused by a familiar pattern or causal mechanism? According to explanatory pluralism, these questions have different answers. Which answer provides the most adequate, efficient, and accurate explanatory information depends on the specific question one wants to answer (or, in other words, on the specific information that is requested in view of the explanatory purposes). Hence, making the specific question as explicit as possible is important for explanatory success.²⁹⁷

Here again, as in the cases of history and social science, and of cognitive science, it appears as though not only is an integral part of motivating explanatory pluralism is the variety of interests and goals one may have in asking an es-question, but additionally it is accepted that what may have appeared to be two inquirers asking the same question (due to uttering the same words in the course of doing so, for instance), upon further analysis, is actually a case of two very different es-questions being asked in accordance with two distinct sets of explanatory interests and aims, which may nevertheless be such that they can be asked in a coarser-grained way by the same utterance. Indeed, it is claimed that there may be several sorts of

²⁹⁷ Vreese *et al.* (2010: 374 – 375)

desired outcomes of es-question within the context of medical science: therapeutic/remedial motivations, prediction, general curiosity and accounting for the unexpected are all taken to be legitimate – yet distinct – aims of inquiry.²⁹⁸

By way of a more concrete example, consider lung cancer. A natural question to ask regarding the condition might be ‘why did person P get lung cancer?’; this question, however, seems on reflection to be interpretable at least as one of the three following specific es-questions:

- (a) Why did person P, who smokes, develop lung cancer, while person P’, who also smokes, did not?
- (b) Why did person P with behaviour B develop lung cancer, while person P’ with behaviour B’ did not?
- (c) Why did person P living in country C develop lung cancer, while person P’ in country C’ did not?²⁹⁹

Such elaborations of different es-questions that appear similar at first glance are reminiscent of the example discussed above of Saul going to the shop. There are differences in the examples, since these three are contrastive questions, whereas the same was not true in the earlier discussion. However, commonalities to the effect that there are different explanatory goals at work in each fully articulated es-question are certainly present.

In those earlier examples, recall that we saw how these different es-questions were such that the preferred eg-answer would vary, and would vary based on the difference between *reference classes* of the es-questions under consideration. Here, it is contended, the es-questions of the lung cancer example ((a) – (c) above) work in much the same manner: whilst attention paid to smoking habit may be taken colloquially to be the answer to the general ‘why did person P get lung cancer?’ question, it is in fact more properly the answer to (b).³⁰⁰ Similarly, knowledge of lung cancer variants might figure most appropriately into an answer to (a), whilst the reference

²⁹⁸ Ibid. (375)

²⁹⁹ Ibid. (375 – 376)

³⁰⁰ Ibid. (376)

class of (c) is perhaps the most fitting present reference class of the three es-questions that may contain matters concerning public health policy.³⁰¹

Once again, then, it appears that what we are left with is an array of explanatory interests, the eg-answers to which are being selected for based upon how well they satisfy the interests of the inquiring agent behind the es-question. Indeed, Vreese *et al.* conclude their discussion of the example with the following telling passage:

*This lung cancer example demonstrates that researchers and practitioners cannot freely select one or the other of the more specific questions to answer the rudimentary question. Instead, something needs to guide them toward their choice. Good scientists think thoroughly about what kind of information they need, and which of the possible more specific explanation-seeking questions will deliver the most adequate, accurate, and efficient explanation in view of their needs. Different epistemic interests will guide medical researchers and practitioners to different kinds of explanation-seeking questions and thus to different kinds of explanations in different explanatory contexts. This means that, in theory, explanatory pluralism holds for the medical sciences.*³⁰²

Here again, the thought does appear to be that it is the different explanatory interests of the inquirer – leading to the advancing of different es-questions – that leads to different kinds of explanations being preferable in different explanatory contexts. Additionally, although the notion that inquirers cannot freely select their es-question may be just as likely to find support from monists (due to the exclusion or preference rules the latter are looking to guide their explanation selection), it seems a fairly reasonable interpretation of the earlier part of the preceding quotation to say that on this view, inquirers form their own es-questions by contemplating which ones will be the most explanatorily adequate, efficient, etc., *in the case in question*. If this is correct, then there can very plausibly not be any general principles regarding explanation, and the explanatory pluralist tenets are satisfied. Moreover, they are satisfied for reasons that do not apply to the subject matter of medical science itself, but rather due to the same variation in

³⁰¹ Ibid.

³⁰² Ibid.

explanatory interests that led to pluralism being motivated in the cases of cognitive science, social science and history. In short, we see that in each of these areas, we can apply a consideration along the lines of the following:

[No] *one explanation has [universal] precedence over another; there is no such thing as the explanation [. . . and the presupposition that there is] does not stand up to critical scrutiny. We must take seriously the context-ladenness of scientific explanation.*³⁰³

In the next section, I shall attempt to provide support for the contention that, at least in the areas philosophy of perception which I am chiefly concerned with, explanatory pluralist methodology of the sort expounded above is capable of being utilised, and will show that the objection to the pluralist approach that contends it is unmotivated is itself unfounded.

3.4: *Explanatory pluralism in the philosophy of perception*

The overall purpose of this chapter has been to show that the utilisation of explanatory pluralism in the philosophy of perception is not doomed to fall prey to the objection that such utilisation is unmotivated. Presently, I am in the position to begin tackling this undertaking. The previous sections provided discussion of numerous different areas of research in which explanatory pluralism is said to be a motivated methodology. Recalling that in each of these areas, the explanatory pluralism in question was motivated by the difference in explanatory interests of the researcher (which led to multiple different sorts of es-questions with different kinds of eg-answers in different explanatory contexts), it seems that what is required to show that pluralism is a motivated view is simply to illustrate that the philosophy of perception also possesses the varying, legitimate sorts of interest that serve to make for the various es-questions and eg-answers. In other words, what must be done in this final section is to establish whether those engaged in philosophy of perception possess the same sort of diverse explanatory

³⁰³ Boylan and O’Gorman (1999: 144)

interests as those engaged in history, social science, cognitive science or medical science. If it turns out that there is a variety of legitimate explanatory interests in the philosophy of perception, then it seems that either the latter is such that explanatory pluralism can be a methodological approach to it which appears well motivated, or else the pluralist position arguably *is* unmotivated, both in the philosophy of perception and in the other disciplines discussed.

It is important to stress that the purpose of this chapter is to establish the disjunctive claim in the preceding sentence. This is to say that the present chapter stops short of establishing that the explanatory pluralist methodology, when applied to the philosophy of perception, actually *does* circumvent the Screening Off Argument in the manner I am envisaging – establishing that conclusion is the task of the remaining chapters. The remainder of this one, meanwhile, is committed to illustrating merely that there is virtually no difference between how the explanatory pluralist methodology can be applied to the philosophy of perception and, as we have seen above, how it has been applied to the social, cognitive and medical sciences. On the basis of establishing this, it becomes possible to submit the overall conclusion that either explanatory pluralism is just as motivated an approach in *all* of these disciplines, unless there is something particular to its application to the philosophy of perception that defeats this, which I shall argue there is not.

The key question for the purposes of this section, then, is whether in working with the relational and representational views of experience, there come times whereby the two theories fall into distinct reference classes of distinct es-questions. To answer this, one should recall how exactly it is that es-questions come to vary to the extent that they do – namely: *difference in explanatory aims and goals of the inquirer*, for it is these that guide the decision-making process as to what the most illuminating (or efficient, etc.) of the es-questions on offer is. It is, I submit, precisely the case that such varied explanatory interests obtain in this case, such that there are distinct

es-questions with distinct reference classes that involve representationalism and relationalism respectively.

Before attempting to argue for this latter point, it will be useful to provide as a 'stepping stone' an example which illustrates how different aspects of perceptual experience may be more relevant in providing eg-answers to es-questions, before trying to compare the respective merits of representationalism and relationalism in providing these eg-answers. If we grant that there are these distinct, legitimate aspects of perception, it can be illustrated that just because focusing on one aspect of perceptual experience is most prudent in providing an eg-answer to an es-question, this does not mean that another aspect of perceptual experience would not be better suited to providing an eg-answer to a different es-question. As we will see in the example that follows, this contention turns out to be true even in cases where it looks as though the different eg-answers are eg-answers to the *same* es-question: whilst there can be a hierarchy of appropriate eg-answers, this is actually established in the example that follows by teasing apart the more specific es-questions to which some respective inquirers *actually* desire some eg-answers (in other words: by focusing upon exactly what it is that they wish to understand), rather than by some general rule that sorts prospective eg-answers into a hierarchy based on the sorts of explanations they provide. This is to say, I wish to show the plausibility of the claim that that of which we are attempting to gain a heightened understanding is what determines where it is most appropriate to look for an eg-answer to the given es-question. Thus, the example will be very similar to that seen in the preceding sections viz. social sciences, psychiatry, and so on, yet applies the same approach seen there to the sorts of es-questions to which philosophers of perception would be likely to seek eg-answers. Note that what follows is an example whereby both relationalism and representationalism may plausibly be utilised: at this point I am merely showing how the explanatory pluralist methodology itself may be applied to the questions that representationalism and relationalism might be appealed to in answering. I leave discovery of stronger examples, where it looks more convincingly as

though only one of the two theories might be of use (often treated in the literature as cases that sway the so-called relationalism versus representationalism debate in favour of one theory or the other), to the following chapter.³⁰⁴

Considering the rudimentary question that we might imagine an inquirer to ask – ‘what is it for S to see *o* as *F*?’ – this is further analysable, in what should hopefully at this point be a familiar way, into more specific es-questions. Which aspect of the initial articulation of this es-question is relevant in each of the more specific versions can be stressed as follows:

- a) What is it for S to see *o* as *F*?
- b) What is it for S to see *o* as *F*?
- c) What is it for S to see *o* as *F*?
- d) What is it for S to see *o* as *F*?
- e) What is it for S to see *o* as *F*?

Here, I take (a) to be unobjectionably the sort of es-question that gets asked in the philosophy of perception. It is, I think, the most natural reading of the rudimentary es-question provided, and is also the one closest to the types of questions with which this thesis is concerned. It is asking for an account of the whole process of a subject visually perceiving an object as instantiating a given property. The second, (b), is on one reading less obviously in *philosophy* of perception territory, but can be read as asking for an account of why a subject is seeing the scene they are, rather than hearing or smelling: it pertains to sensory modalities. Insofar as some work in the discipline focuses on sensory modalities, I do not think treating this as an es-question arising in the philosophy of perception is too great a leap. Versions (c) and (d) can be read as pertaining respectively to the particularity of objects perceived and to how we predicate properties to objects we perceive. I discuss both cases in the following chapter. Finally, (e) is a contrastive

³⁰⁴ These stronger cases, though more useful to my ultimate aim, require argument in support of the contention that they qualify as such. Though necessary to my overall conclusion, I deem such arguments extraneous to the point currently at issue – which is to illustrate the application of the methodology of explanatory pluralism itself – and so delay their discussion for now.

question seeking an account of *which* property is predicated to a perceived object.

For clarity, here is the list of more specific es-questions that fall out of these readings:

- a) What is it for S to see *o* as *F*?
- b) What is it for S to see – rather than hear, smell, or taste – *o* as *F*?
- c) What is it for S to see *o* – rather than *o** – as *F*?
- d) What is it for S to predicate the property *F* to object *o* in their visual experience?
- e) What is it for S to see *o* as *F*, rather than seeing it as *G*?

When analysing the motivation of explanatory pluralism in previous sections, one notion that was appealed to was that of a reference class: the range of options salient to providing the eg-answer to the es-question. This was useful in relation to contrastive es-questions, or those whose eg-answers had causal explanations. The difference in reference class, it was held, served as a plausible indicator of difference in explanatory interests. Since the latter are what is required to motivate explanatory pluralism, I am arguing, it will be useful to consider this notion here.

Examining es-questions (a) – (e), we see that three of them are contrastive, whilst two – (a) and (d) – are not. I return to these two below. Of the remaining three, identifying the reference class of each is straightforward: the range of sensory modalities serves as the reference class for (b); the range of mind-independent objects serves as that of (c); and the range of predicable properties that can possibly be instantiated by mind-independent objects serves as that of (e). These are all different reference classes, and serve to highlight the different explanatory interests realised by the respective inquirers of (b), (c) and (e).

Finally, (a) and (d) are not contrastive, but instead seem interested in identifying the nature of something. In the case of (a), it seems that what is sought is an identification of the nature of the perceptual experience as a whole. Since the eg-answer to this is such that it could possibly contain the eg-answer to (b) – (d) as constituent parts, I take it to be plausible that the

explanatory interests realised by (a)'s inquirer differ from those realised by the inquirers behind the other es-questions. Lastly, (d) is also looking to get at the precise nature of something in a non-contrastive manner.³⁰⁵ In this regard it is similar to (a), but the predication of properties in visual experience seems a constituent part of what it is to have a visual experience, in the sense that one can intelligibly ask about it and view as explanatorily irrelevant many other aspects of what it is to have a visual experience. Conversely, one could ask (a), be given the perfect eg-answer to (d), and view this as an incomplete explanation.

This overall point generalises: an explanation that perfectly answers (a) will have as its reference class the range of sensory modalities, but this will serve very poorly as a reference class for (e), which seeks an eg-answer relating to which property is predicated to a visually perceived object. For the latter question, the visual modality is held fixed, and so matters of other sensory modalities are not relevant to providing a satisfactory eg-answer. More generally: certain aspects of perceptual experience are simply less relevant than others in arriving at a satisfactory eg-answer to legitimate es-questions. Moreover, it is the es-question *itself* – and the explanatory interests which generate it in the first place – that determine what is relevant. As Bouwel puts it:

*By making the different possible explanatory requests explicit [. . .] the motivation and the explanatory information required will be taken into account. It can be shown that one [. . .] fact can be the subject of different questions, and hence of different forms of explanation. Consequently, taking into account the explanatory question is not something of secondary importance, as it decides on which form of explanation will be used. To be able to answer these different kinds of explanatory questions we will need different forms of explanation.*³⁰⁶

³⁰⁵ At least not in a manner that contrasts with anything in particular, as is the case for the other es-questions that possess an obvious reference class. The case could be made that the inquirer is interested in how predication of properties to perceived objects contrasts with *everything* else, but this is not an important point.

³⁰⁶ Bouwel (2004: 308)

If the above analysis is correct, then it appears that different aspects of perception can generate different es-questions due to divergent explanatory interests involved in investigating them, and these es-questions are such that the way to the most appropriate eg-answer will differ. It has also been shown that at a sufficiently coarse-grained level, these es-questions can all appear identical, despite actually being distinct. Importantly, these differences map onto the explanatory interests of the inquirer, which is why there are no generalised rules of explanatory exclusion or preference at the global level. Yet this is precisely what was required in order to satisfy the claim that explanatory pluralism was motivated in researching cognitive science, medical science, history, etc., hence it seems that in the absence of any subject-specific defeaters present in the philosophy of perception – which are not obviously present – we must conclude that explanatory pluralism is at least as motivated in this area as in the others examined earlier.³⁰⁷ Of course, this allows equally for the view that explanatory pluralism is motivated in *none* of those research areas mentioned, but this is a hugely strong monist claim, and without substantive argument offered by the monist, it is hard to see why we ought to abandon what is potentially a highly intellectually lucrative approach to such a diverse range of quests for understanding.

Perhaps one relevant difference in terms of types of explanations is that in the previous sections where explanatory pluralism was held to be motivated in other disciplines, the examples used involved *causal* explanations. Since the particular question in the philosophy of perception that I am concerned with is about the nature of perception, it might be held that this difference precludes inference from the methodology's motivation in those cases to its motivation in the philosophy of perception.

³⁰⁷ By subject-specific defeaters, I mean something particular about engaging in the philosophy of perception that undercuts my argument supporting the view that pluralism is motivated. Though I do merely say above that the presence of such a defeater is not obvious, perhaps the most obvious potential candidate would simply be that perception is fundamentally just not an area of research that admits of more than one kind of explanation, but this would be question-begging.

I think the above example es-questions illustrate that this is not a relevant difference that outright blocks the conclusion that explanatory pluralism could be motivated in the philosophy of perception. Certainly, it does not obviously block the conclusion that it is *as* motivated in the philosophy of perception as it is in the other disciplines considered. More than one of the es-questions considered did not involve a causal explanation, yet were in the philosophy of perception remit, and upheld the broader point that an array of explanatory interests could be shown to operate across the es-questions considered. This can be argued, just as in preceding sections, to show that different eg-answers are more suitable for addressing different es-questions.

If there is any further reason why this causal/non-causal distinction renders application of explanatory pluralism unmotivated in the philosophy of perception, it is not clear to me what it is. I have shown that the methodology broadly holds up in dealing with the sorts of es-questions occurring in the philosophy of perception, which do not seek causal explanations. Consequently, I take the onus to be on the party voicing continued scepticism towards explanatory pluralism's motivation in this discipline to argue the point.

I have argued that what is required to motivate explanatory pluralism is simply the existence of different explanatory interests, such that these result in different es-questions being generated which themselves have different – more or less fitting – eg-answers, to the effect that the preferred possible explanation depends upon the explanatory interests of the inquirer. I have supported this claim by examining how explanatory pluralism is supposed to be motivated in other disciplines – namely: history and social science, cognitive science, and medical science. In all of these, the claim I am arguing for appeared to hold, for different explanatory interests *do* appear to lead to distinct es-questions in each case. Thereafter, it was demonstrated that the same can be said for the philosophy of perception, and there are no obvious countervailing reasons for nevertheless rejecting the view that explanatory pluralism might be motivated in the latter discipline if it is motivated at all in any of the others. If the foregoing analysis is correct, then it appears the

utilisation of explanatory pluralism in the philosophy of perception cannot be blocked by way of objecting that such an approach is unmotivated. This is important because, as we shall see throughout the remaining chapters, utilisation of explanatory pluralism allows an account to be developed which successfully overcomes the Screening Off Argument, and which entails the compatibility of relationalism and representationalism.

4: *The distinct explanatory interests of relationalism and representationalism*

I identified the Screening Off Argument in §0.3 as a frontrunning reason for maintaining that relationalism and representationalism are incompatible, and therefore as the argument I wish to overcome. The first part of the approach I have in mind for overcoming the Screening Off Argument can be realised based on the concepts introduced in the previous chapter. In particular, it will rely on the notion of explanations being reducible to pairings of es-questions and eg-answers. Broadly, the thought is that by focusing on certain es-questions addressed by relationalism and representationalism respectively, it is possible to identify explanatory interests that will only be served by one theory or the other. Specifically, this will be the case when there can be identified an es-question which is such that an eg-answer is provided to it by only *one* of relationalism or representationalism, but *not both*. It is identifying es-questions such as this, and arguing that they conform to this constraint, that will be the task of this chapter. The role that such es-questions play in my attempt to reconcile relationalism and representationalism will become clear in the following chapters.

4.1: *The interests of relationalism*

Starting with relationalism, I now commence in an effort to identify an es-question which is such that an eg-answer is offered to it by relationalism but not by representationalism. There are two possible ways I see of accomplishing this. The first is to identify an es-question which is *de facto* in such a position, which is to say that there is currently no eg-answer offered to it on behalf of representationalism. The problem with such an approach is

that there is nothing to stop the representationalist from thereafter formulating and offering an eg-answer to this question, thereby undermining its support of my overall argument. Conversely, the second way to accomplish this task is to identify an es-question which is such that representationalism cannot, due to its commitments, provide an eg-answer to it. This es-question, if it really is that way, does not allow for a representationalist eg-answer being offered at any time, and thereby circumvents the worry of the first method outlined above. If possible, therefore, es-questions that satisfy the constraints imposed by the second method are the es-questions that I shall be attempting to identify in this section.

I take it that a fruitful way of identifying such an es-question would be to consider just what it is that a relationalist theory does that a representationalist theory does not. One way of articulating this that is agreeable both to a reconciliatory as well as austere view is that relationalism fundamentally explains aspects of perception by appealing to the mind-independent objects of perception themselves, as well as the circumstances of the experience (point of view, lighting conditions, and so on). With reference to these, the relationalist can articulate a particular relation of direct acquaintance that the subject stands in with the objects that they perceive in their environment. This acquaintance relation is then used to explain aspects of perceptual experience in need of explanation. This articulation is acceptable to both reconciliatory and austere theorists because it is compatible both with the claim that representation occurs in perceptual experience (reconciliatory) and with the claim that *all* aspects of perceptual experience are to be explained by the subject's acquaintance with the mind-independent objects they perceive.

The two sorts of theorist will certainly disagree over the latter of those claims being true. It is in the nature of austere relationalism to claim that the acquaintance relation, and nothing more, explains all that is required of a philosophical theory of perception; a reconciliatory view, by its nature, does not accept this claim. This is compatible with both a relation of acquaintance

and representation occurring during our perceptual experience. Where the two sorts of view disagree is in which of these elements does the heavy lifting in explaining perceptual experiences.

Relationalism's explanatory focus on the relation of direct acquaintance, I submit, tells us about something that relates to the es-question for which we are presently looking. In particular, it tells us what its eg-answer will look like. The eg-answer will appeal to the direct acquaintance relation - itself cashed out in terms of mind-independent physical objects themselves, as well as the point of view, lighting conditions, and so on, from which they are perceived. These are what fall within the explanatory purview of relationalism exclusively, and so if the es-question of the sort to be identified *does* exist, the criterion that relationalism alone can provide an eg-answer to it restricts its eg-answer to this domain. This gives us a concrete starting point.

From here, the next move would be to identify the counterpart es-question to this form of eg-answer. This may strike some as a manner of formulating the es-question that is undertaken for the *ad hoc* reason of establishing a reconciliatory theory of perception. Against this charge, I counter that if such an es-question can (by any means) be found, along with a comparable es-question to which only representationalism can provide a satisfactory eg-answer, then a reconciliatory view becomes overwhelmingly plausible. It is simply the case, I argue, that both austere versions of relationalism and representationalism turn out to be wrong in - and because of - instances such as these. I take the consequences for the relationalism versus representationalism debate, which I will argue these instances have, to overcome any allegations that the approach is *ad hoc*.

A reconciliatory view being correct is one plausible outcome of this. Another available option is that an entirely distinct view may explain everything we wish to explain. Consequently, it is not strictly speaking *inevitable* that a

reconciliatory view falls out of the discovery of these es-questions.³⁰⁸ A wholly distinct view would not only need to be just that, but would also have to account for the appeal that both representationalism and relationalism have, whilst *also* accounting for why they were both ultimately mistaken. I take this to be a less plausible option than one that can account for all of this without also being required to differ substantially from relationalism and representationalism.

Let us return to identifying our es-question. We now know that its eg-answer must appeal ultimately to the mind-independent objects themselves, and the salient circumstances (point of view, lighting conditions, and so on) within which they are perceived. To progress this further we might consider, then, what role these play. I suspect, though, that this would make for a quite a long and complex answer. It also seems ill suited as it stands to serve as our target es-question: it is sufficiently generic and has a wide enough scope that the representationalist would likely be able to provide an eg-answer to it. Nevertheless, the overall approach of analysing what work the appeal to mind-independent objects and the circumstances in which they are perceived does for the relationalist account seem a plausible approach to take. The task will be confounded by the need to narrow the scope and suggest a more specific es-question, but this will not render it insurmountable.

The upshot of this thought, I suggest, is that a way of sharpening our focus towards what we are looking for is to specify the *sort* of role that mind-independent objects and the circumstances within which they are perceived play for relationalism, and to request an explanation of *that* role. The thought is that perhaps they play a role within the relationalist explanation of some aspect of perceptual experience which is intelligible on relationalism alone. That is: the mind-independent objects and circumstances of

³⁰⁸ Though of course one could restrict the context of discussion to include only promising views widely discussed in contemporary philosophy of perception literature, which would have the effect of shifting far closer to inevitability.

perception could *not* play this role on a representationalist picture, whatever other roles they might feasibly play for that theory.

Broad candidates for different sorts of roles can be found by considering certain aspects of perceptual experience – namely: phenomenal, behavioural, and epistemological.³⁰⁹ The candidate I have in mind is the latter, as I take relationalists to generally consider their view to provide a satisfactory account of the epistemological aspects of perceptual experience.³¹⁰ Their account enjoys intuitive plausibility, because a naturally appealing explanation for how we can know about the objects in our environment via perception is that perception involves being directly related to those mind-independent objects themselves.

The means of approaching the target es-question, then, will be to consider what *epistemological* role is played by the apparatus that is central to the story told by relationalism. In order to do this, we should first consider the positive epistemological story told by relationalism.

4.1.1: *Conclusive reasons*

In Chapter 2, I set out the commitments of relationalism for visual perception. On Brewer's articulation of relationalism, this involved accounting for seeing in terms of propositions of the form '*o* looks *F* to *S*', where *o* is the mind-independent object perceived, '*F*' is some property, and *S* is the perceiving subject. The idea was that *o* looks *F* to *S* iff (i) *S* stands in a relation of direct acquaintance *R* to *o*, from a standpoint *c* – which is such that *o* bears visually relevant similarities to paradigm exemplars of *F* when viewed *from c* – and (ii) *S* registers that *o* bears visually relevant similarities to paradigm exemplars of *F* when viewed from *c*. According to relationalism, if this captures something of the basic structure of perceptual experience,

³⁰⁹ See Logue (2014)

³¹⁰ Ch. 6 of Brewer (2011) is dedicated to highlighting precisely this, and Martin (2004, 2006) famously accounts for hallucination in a wholly epistemic way.

then any positive account of perceptual knowledge will need to adhere to it.³¹¹

In the most straightforward cases of knowledge, there is reason to suppose that factual knowledge is dependent upon possession of concepts relevant to the propositional articulation of that knowledge.³¹² One cannot know what one lacks the appropriate concepts to express; this is not to claim that one must be capable of articulating the proposition in question – but one must possess the concepts required to do so. The relationalist can say that, provided the subject possesses the relevant concepts and is acquainted with the epistemic grounds that render their perceptually based judgement correct (in a manner to be discussed below), conscious acquaintance with an object “normally makes application of ‘*F*’ in judgement *evidently* correct”, as in such cases the subject possesses the relevant concepts and is viewing the object from an appropriate position and in correct conditions.³¹³ Hence, seeing the object constitutes the subject’s *reason* (in the sense of providing epistemic grounding) for judging that *o* is *F* – i.e. “experience acquaints us with the grounds for empirical truth.”³¹⁴ This can be expressed in the relationalist framework from earlier by stating that, on the supposition that *o* is *F*, “*o* makes an application of ‘*F*’ correct: *o* itself is what makes ‘*o* is *F*’ true and in this sense constitutes a reason to apply the predicate.”³¹⁵

The representationalist can rightly say the same, and would have an equally legitimate claim that the object *o* actually being *F* is what renders *S*’s judgement that *o* is *F* is correct. What I shall argue in what follows is that the representationalist cannot place the same emphasis on *reasons* as the relationalist can. To illustrate the point, a question to consider is: what are the respective epistemological contributions of (a) the mind-independent

³¹¹ I borrow some of Brewer’s terminology in what follows simply because it renders the point to be made clearly and precisely. Not all relationalists would spell things out in exactly this fashion, but differences in precise articulation do not affect the broader point being made.

³¹² *Ibid.* (142)

³¹³ *Ibid.* (143)

³¹⁴ *Ibid.*

³¹⁵ *Ibid.* (142)

object perceived, and (b) the standpoint from which the object's being a particular way is made apparent to the subject?³¹⁶

I shall state the relationalist's response to each of these in turn, before analysing them. In response to (a), Brewer states:

(A) *It is a necessary condition upon a perceiver's having an experientially based reason of the kind that I am elucidating here to apply the concept 'F' in judgement that she should be consciously acquainted with what is in fact a reason for such application, namely, a direct object of perception, o, that is in fact F. That object o itself is a reason in the relevant sense to make the concept application in question in judgement.*³¹⁷

On this view, then, the object itself plays the role of providing a *reason* for the subject to predicate a property to that object. More precisely, it provides the subject with a reason to judge that their environment is a certain way (that property *F* is instantiated by object *o*, say) that is *conclusive* – that is: a reason that does not leave the possibility open that the judgement is mistaken.³¹⁸ This is to be taken as in Chapter 2, where McDowell's epistemological disjunctivism was discussed. That view holds that there is an epistemic asymmetry between successful (veridical) perceptions and corresponding (subjectively indistinguishable) hallucinations. In the so-called 'good case' of successful perception, one's experience provides one with warrant to make a judgement about their environment that is conclusive in the sense specified. Conversely, the so-called 'bad case' does provide the subject with warrant for making such a judgement, but in this case warrant that is compatible with the judgement being mistaken. The view takes a subject's experience of things appearing a certain way to constitute a *disjunctive fact*: either an objective state of affairs becomes manifest to the subject, or else a situation in which it is, to the subject, *as though* an objective state of affairs becomes manifest to them (when that is not how things are).

³¹⁶ This question is adapted from one in Brewer (2011: 147), credited to Anil Gupta.

³¹⁷ *Ibid.* (147 – 148)

³¹⁸ Soteriou (2016: 119)

Importantly for our present purposes, the McDowellian position has consequences for our perceptual knowledge – particularly: it impacts the epistemic status of the *reasons* a subject has for judging their environment to be a certain way. I alluded above to the ability both relationalists and representationalists have to appeal to mind-independent objects as constituting reasons *why* the subject is correct in judging their environment to be thus and so. Distinguishing reasons why such a judgement may be correct from the epistemic reasons (or grounds) upon which the subject might base that judgement, however, highlights an important point. In considering epistemic grounds for a subject’s perceptually based judgements, epistemological disjunctivism allows its proponent to simultaneously hold that these grounds stop short of the fact in the bad case, but not in the good case. We also saw that epistemological disjunctivism does not entail metaphysical disjunctivism, so the success of the former cannot be said to hang on the success of the latter.

With all of this in mind, we can unpack (A) above. Brewer’s claim is that to have the sort of epistemic grounds that do not fall short of the fact, one must be directly acquainted with the mind-independent object and its properties. The thought is that the object itself (and its properties) are the only grounds for the subject’s judgement that their environment is a certain way that are incompatible with that judgement being mistaken. Mere appearance (that falls short of the fact, in the relevant sense) is compatible with the judgement being mistaken, and so does *not* constitute the sort of reason envisaged.

Meanwhile, (b) seeks to establish the epistemological role of the subject’s standpoint in making conclusive reasons for perceptual judgement apparent to them. Specifically, it asks what epistemological bearing the circumstances of perception (*C*) have on the subject making their perceptual judgement. Since *C* itself details the way the subject becomes acquainted with *o* (and thus acquainted in the good case with conclusive reason for their perceptual judgement), the answer to (b) had best establish how such a reason can be realised by a subject. One way in which (b) has been responded to is in terms

of how mind-independent objects perceptually look to a subject who is directly acquainted with them. One account of this discussed in Chapter 2 cashed it out in terms of *o* bearing visually relevant similarities to paradigm exemplars of *F*. Though not the only way of providing an answer to (b), Brewer offers the following:

(B) *It is a further necessary condition on that very reason coming to light in her experience that she be acquainted with o from a point of view and in circumstances that enable her registration of the appropriate visually relevant similarities that it has from that point of view and in those circumstances with the paradigms that are involved in her grasp of the concept 'F'.³¹⁹*

The point of view and circumstances, *C* – via which Brewer’s account claims that *o*’s visually relevant similarities to paradigm exemplars of *F* become apparent – play a central role in the subject appreciating the reason articulated in (a) above. More precisely: it is a necessary condition on *S* grasping this *conclusive* reason for applying *F* to *o*, that they are acquainted with *o* from some standpoint *C**.

Satisfaction of **(A)** without **(B)** renders possible cases where *S* is acquainted with what is in fact a reason to judge that *o* is *F*, yet this fails to be made evident to *S*.³²⁰ Satisfaction of **(B)** without **(A)**, meanwhile, allows illusory experience of an object that is *G* but not *F*, and is viewed in conditions which cause it to have visually relevant similarities with paradigm exemplars of *F*, leading to cases of understandable errors in judgement based on the experience.³²¹ Hence, establishing both **(A)** and **(B)** precludes respectively that a subject has equally strong justification for their perceptual judgement in both the good and bad cases, and that the subject has legitimate grounds for doubt in the good case. Looked at one way, these are both strengths of the position; though one’s *modus ponens* being another’s *modus tollens*, such preclusions can also be looked at as the bases for two objections to the view.

³¹⁹ Brewer (2011: 148)

³²⁰ Ibid.

³²¹ Ibid.

In the first case, the objector argues that it is counterintuitive to say that the strength of a subject's justification across the good and bad cases varies – we should want to say that the hallucinating subject is just as justified in their judgement as the perceiving subject *just because* the respective experiences are subjectively indistinguishable.³²² The idea is that, given the nature of the subject's experience in each case, and the introspective indistinguishability across cases, the epistemically appropriate response to each experience is the subject making the same judgement, with identical credence in each case.³²³ However, one can uphold that there is a match of sorts in epistemically permissible responses to cases whilst denying that there is epistemic symmetry in terms of how conclusive the subject's reasons for perceptual judgements are across good and bad cases.³²⁴ Hence what is going on in non-veridical experience is to be explained *derivatively* on the basis of what is going on in veridical experience: in the latter, one has access to conclusive warrant for one's perceptual judgement; in the former, one only seems to have such access, and is not in a position to tell that one does not.³²⁵

In the second case, the objector seeks to argue that it is plausible that subjects can have legitimate grounds for doubting their perceptual judgements in good cases: for instance, testimony of others may provide the veridical perceiver with grounds for doubt that are legitimate, even though they are mistaken.³²⁶ However, here the epistemological disjunctivist can appeal to the potential for the double dissociation between **(A)** and **(B)** – it is consistent with the subject having a conclusive reason to perceptually judge that *o* is *F*, for instance, that they fail to appreciate that they have this reason. Consequently, the subject in the good case can be sensitive to purported counterevidence to their perceptual judgement, even if that purported counterevidence is mistaken.³²⁷ As Soteriou puts it, these cases are

³²² Soteriou (2016: 136 – 137)

³²³ *Ibid.* (138)

³²⁴ *Ibid.* (139)

³²⁵ *Ibid.*

³²⁶ *Ibid.* (140)

³²⁷ *Ibid.* (141)

such that “you are deprived of the awareness that your experience puts you in a position to know things about your environment.”³²⁸

Having established, in line with (A) and (B) that both foregoing considerations constitute necessary conditions for perceptual knowledge, it follows that it is not enough that S is acquainted with relevant reason for judging *o* to be *F* in order for S to *know* that *o* is *F* – they must also appreciate the reason in question. In other words, they must be aware of *and* accept the reason they have to apply the relevant concept(s) to the object(s) in their surroundings.

Brewer cashes out this acceptance of the sort of reason in question in terms of *conceptual registration*. Conceptual registration constitutes the subject’s appreciation of the reason in question *as being* the conclusive epistemic reason for their judgement that it is. The subject acquainted with conclusive reasons for perceptual judgement undergoes conceptual registration, accepting – in a manner constitutive of knowledge – that their environment is a certain way. As Brewer puts it: “acquaintance has to be combined with conceptual registration and endorsement for the acquisition of knowledge.”³²⁹

We now have a clear articulation of the respective roles played by *o* and *C* in the acquisition of perceptual knowledge – thereby answering questions (a) and (b) above of the epistemological roles played by each. It will now therefore be prudent to assess whether either of these allow us to formulate the target es-question. I do not think that the analysis in (B), pertaining to the circumstances of perception and point of view, are such that they will do the job. This is because one of the criteria of our target es-question is that representationalism is unable to provide an eg-answer to it. As we saw, the epistemic contribution of the circumstances of perception and point of view are closely bound up with conceptual registration – the notion that S must endorse, as well as be acquainted with, reasons for applying the property *F*

³²⁸ Ibid.

³²⁹ (2011: 145)

to the object *o*. Conceptual registration, as drawn up here, sounds a great deal like the representationalist's conception of representation. For instance, conceptual registration as taken to be "an irreducibly primitive notion that [can be conveyed] by appeal to such familiar commonsense activities as noticing, recognising, *and seeing as*."³³⁰ The emphasis in the preceding quotation echoes closely representation as I have defined it in Chapter 1. Thus, I think it arguable that the functional role played by each notion is similar enough that an es-question based on the deliverances of **(B)** would plausibly allow for a representationalist eg-answer.

Giving credit where it's due, Brewer anticipates this objection, contending that two considerations support a distinction between representation and conceptual registration:

*First, acquaintance itself does not presuppose conceptual registration. Hence registration is not in this sense part of the fundamental nature of our basic perceptual relation with the mind-independent physical world. Second, registration is something that subjects themselves do [. . .] Hence, again, this is not a feature of our most basic deliverances of our perceptual systems in our conscious experiences of the world around us.*³³¹

For reasons that are worth pointing out for my later purposes, I think this statement stops short of showing that representation fails to play a central explanatory role in our perceptual epistemology.

The first claim – that acquaintance does not presuppose conceptual registration – highlights a difference in importance placed upon conceptual registration and representation respectively for Brewer's view and for representationalism. The thought is that since representation occupies a position of foundational importance for representationalism, whilst the same cannot be said for conceptual registration's place in Brewer's relationalism, the two are not identical. This is fine, as far as it goes. It does not preclude a representationalist answering an es-question about the

³³⁰ Ibid.; emphasis mine.

³³¹ Ibid. (145 – 146)

epistemological role circumstances of perception *C* play in perceptual knowledge, however. If one assumes that what I am calling austere theories of perception are the only viable option, then this difference would be sufficient. Reconciliatory views can allow for the same thing (say, a subject's capacity for perceptual representation) to occupy explanatory roles of varying importance depending on what is to be explained, however. Consequently, one cannot claim that conceptual registration is obviously distinct from representation in the manner required on the basis that each respective theory assigns them varying levels of importance.

To the second claim: that subjects themselves carry out conceptual registration clearly entails no difference between that process and representation as I am understanding the term. Promising theories of representation can happily argue that the subject carrying out the process autorepresentation *is* in fact a most basic feature of our perceptual experience, and so I do not accept this as a reason for drawing a distinction between the two. That said, if it turns out that I am wrong about either of these things, then all this means is that the deliverances of **(B)** might more plausibly lead to an articulation of the target es-question. Consequently, I shall say no more on the matter, as I think the deliverances of **(A)** are much more clearly suited to that purpose.

Recall that **(A)** involved the object *o* being *F* itself constituting a reason for the subject judging *o* being *F*. More generally, one might say that mind-independent objects themselves constitute *reasons* for the correctness of applying concepts (of the sort that **(A)** establishes as a necessary condition on perceptual knowledge) in judgements about those objects.³³² In other words, one might say that *o* itself being *F* constitutes a *conclusive reason* for the correctness of *S* judging *o* to be *F*.

It was stipulated above that the eg-answer to our target es-question is restricted by the criterion that only relationalism, and not representationalism, can provide it. In real terms, this precludes the subject

³³² Ibid. (156)

being in a representational state – i.e. *S* representing *o* as being *F* – from doing the explanatory heavy lifting. I take the notion of the object constituting *reasons* for the correctness of the subject’s application of properties to that object as satisfying this constraint, so now turn to articulating the es-question I have in mind, and then supporting this contention.

Since *o* in fact being *F* is a *conclusive reason* for *S* to judge that *o* is *F* – and this in turn constitutes an appeal to mind-independent facts about the object itself, as opposed to *S* being in a representational state compatible with this judgement being mistaken – I take simple identification of these conclusive reasons to appeal to relationalist but not representationalist apparatus. Consequently, an articulation of the target es-question I have in mind may be the following:

(REL-1): How does experience provide the perceiving subject with *conclusive* reasons for applying certain empirical concepts in judging their environment as being a particular way?

The ‘certain’ empirical concepts here just are those concepts that the subject is required to possess and deploy to judge the environment in question as being a particular way – that ‘particular way’, in turn, just being the way the environment actually is. As we have seen, the relationalist provides an eg-answer to (REL-1) by referring to the mind-independent objects themselves, together with their properties, and the relation of direct acquaintance that the subject stands in to these. The representationalist, in contrast, cannot appeal to the subject representing their environment as being a certain way – even the correct way – as this is compatible with the falsity of their judgements, on the widely accepted position that representation requires the possibility of misrepresentation. I take it that (REL-1) itself is an acceptable articulation of a fair question that one might expect a philosophical theory of perception to answer. It seems therefore that the representationalist cannot dispute the legitimacy of (REL-1) as an es-question.

Perhaps the representationalist will appeal to a representation that is *caused* by the mind-independent objects and properties that the relationalist appeals to, but such an appeal would strip any of the representational apparatus of any of its explanatory power. An eg-answer to (REL-1) must show how experience provides the subject with reasons for judgment that are conclusive, and according to this purported eg-answer what is accomplishing this is the relation of causation that the representation stands in to mind-independent objects. Once again, the mind-independent objects and relations to them are providing any success enjoyed by the eg-answer, and representation itself is not contributing anything of explanatory value to the situation.

At this point, the only way out I can see for the representationalist is to bite the bullet, and deny that experience provides us with conclusive reasons for applying certain empirical concepts in judging our environment as being a particular way. This approach has several effects. One of these, happily for the representationalist, is that representationalism is no longer obliged to provide an eg-answer to (REL-1). Instead, what they arguably must now do is provide an error theory, explaining why it is a mistake to hold that experience plays this role of providing us with the relevant conclusive reasons, as well as why some have been prone to commit this misstep. Additionally, they must either provide an account of what does provide us with these conclusive reasons or adopt the sceptical position of denying that we have them altogether. Whatever else this approach might accomplish, it devalues (REL-1) for the representationalist. This is the important consequence for my purposes, as it renders (REL-1) an es-question which is such that only one of relationalism or representationalism - in this case the former - can provide an eg-answer to it.

My overall argument to be made is the following: provided that it is possible to identify es-questions like this, with at least one pair such that only representationalism can provide an eg-answer to one es-question and only relationalism can provide an eg-answer to the other, then it is arguable that our philosophical theory of perception should involve both relationalism and

representationalism. I will introduce the other apparatus required for making this argument in Chapter 5, but it is prudent at this evidence-gathering stage of the argument to note and dismiss one possible means of resisting it.

We have seen that one approach to being confronted with es-questions of the sort I am looking for is to devalue them: to regard them as questions that our philosophical theory of perception ought not to answer. Whilst this would count against my appeal to any of the es-questions one would be willing to devalue in this way, it would come at a price. Once a plurality of such es-questions have been identified, devaluing all or even most of them would entail that our philosophical theory of perception ought to explain far less than would seem obvious. The price, then, would amount to “devaluing the very debate one is engaged in.”³³³

Finally, it is worth noting that, on another way of viewing this situation, it can be held that if such conclusive reasons are necessary for knowledge, then the alternative es-question, ‘how does perception make immediate knowledge of one’s environment possible?’, looks like it would serve a comparable role to (REL-1). In either case, the representationalist needs to explain away rather than explain conclusive reasons of the sort envisaged, as their view does not provide any illuminating account of them. I have chosen for (REL-1) the approach involving fewer assumptions, however I do think that the alternative suggested here is even more obviously an es-question to which our theory of perception must provide an eg-answer.

4.1.2: Demonstrative reference

I have argued in Chapter 2 (especially in §2.3.1) that relationalism is equipped to accommodate our capacity for demonstrative reference in a manner that representationalism simply is not. I shall now identify another

³³³ Logue (2014: 236)

es-question to which relationalism but not representationalism can provide an eg-answer, before providing the reader with some reminders of the considerations discussed earlier that have led to this conclusion. The es-question I have in mind is, straightforwardly, the following:

(REL-2): What makes demonstrative reference possible?

A common thread running between the considerations from Chapter 2 that support only relationalism being able to provide an eg-answer to this involves two claims. The first is that demonstrative reference involves a particularly close connection to the mind-independent object referred to, and the second is that this close connection does not find a natural home in the story told by representationalism, yet can be readily accommodated by relationalism. I shall now address these claims.

Support for the idea that demonstrative reference involves a close connection to the mind-independent object(s) referred to was seen when we examined Jackson's case of Mary the colour scientist. Recall that this case involved the brilliant scientist, who knew everything that there was to know through testimony about colour and colour experience, yet had lived all her life in a monochrome laboratory without access to any phenomenal colour experience of her own. Upon leaving the room and seeing colour for herself for the first time, the question posed was whether or not this constituted Mary learning anything.

This knowledge would constitute knowledge that could not be derived from testimony alone, but that could only come from experience. As we saw, the relationalist considers that Mary does learn something: when she sees the colour red for the first time, she is *acquainted* with the colour red for the first time, and what it is like to see that colour. Mary learns what it is like to see the colour red *because* she becomes newly acquainted with that colour and what it *itself* is like (i.e. becomes acquainted with its intrinsic properties). If this knowledge about red things is afforded her uniquely via her new experience, and not available to her via testimony, there is an argument to be made that we might expect the content of the knowledge Mary gains to

involve a close connection between red things, on the one hand, and the phenomenal character of her experience of redness, on the other.³³⁴ This close connection, it seems, can be captured (perhaps only) by means of a demonstrative: Mary now knows that instances of redness are like *that* (demonstrating the phenomenal character of her experience).³³⁵

Some insight into what this contrast in knowledge gained by Mary when she sees colour for the first time is provided by Campbell, who writes:

*The contrast between the knowledge you have now, on the basis of a look at the objects, and the knowledge you had before of the existence of objects with particular functional roles, is that when you see the thing, you are confronted by the individual substance itself. On seeing it, you no longer have knowledge of the object merely as the postulated occupant of a particular functional role. Your experience of the object, when you see it, provides you with knowledge of the categorical grounds of the collections of dispositions you had earlier postulated.*³³⁶

Upon attending to the actual red objects, then, Mary came to know not only her exhaustive body of testimonial knowledge of redness – functional roles occupied by the property and things that instantiate it, say – but she came to be newly acquainted with the objects that occupy that functional role.

This acquaintance, it can be argued, is the kind of connection that can be captured only by means of a demonstrative, for it is with reference to ‘*that* red thing’, rather than ‘the red thing occupying functional role *x*’ that Mary might newly express about red things. Nevertheless, it is also a result of Mary referring to her own experience, in particular to its phenomenal character. On the relationalist picture of phenomenal character, to refer to phenomenal character of an experience is to refer to the mind-independent objects that it is an experience of, for the latter constitute the former.³³⁷

If what Mary newly learns is accessible to her via the phenomenal character of her experience, and the content of the knowledge involves a close

³³⁴ See Logue (2012: 229)

³³⁵ Ibid.

³³⁶ Campbell (2002: 115 – 116)

³³⁷ Ibid. (116)

connection between that phenomenal character and the object experienced capturable only by means of a demonstrative, then we can see how relationalism accommodates this so well: the view holds that the referent of the demonstrative reference is precisely what constitutes the phenomenal character of the experience that enables her to refer to the redness demonstratively.

Having established that relationalism can accommodate the close connection to mind-independent objects and experience that demonstrative reference involves, let us consider once more the prospect of representationalism for accommodating this connection. If it cannot do this, I submit, then the view cannot provide an account of how demonstrative reference is possible, and so (REL-2) will go through as another es-question to which relationalism but not representationalism can provide an eg-answer.

We saw in Chapter 2 the example of representationalism accounting for my experience of a white cup. The example comes from Campbell, who notes that if you and I are sitting in exactly similar prison cells looking at exactly similar scenes, and each thinking ‘that cup is empty’, then the content of our experience is identical; further, since we are by hypothesis looking at different cells, scenes and cups – i.e. at different particular objects – neither of our perceptions are representing the *particular* things we are looking at.³³⁸

This squares poorly with the thought that when we each think ‘that cup is empty’ we both seem to know which particular cup we are talking about; for all that has been said so far, representationalism has not yet provided an explanation for our possession of this knowledge. We have seen that representationalist efforts to build the capacity for demonstrative reference into the story they tell about representational content have resulted in accounts where experience of objects pre-supposes our ability to think about them, and therefore cannot be appealed to in explaining that capacity.

³³⁸ Ibid. (124)

Views that fall prey to this objection include those which attempt to give representational contents a ‘demonstrative element’, whereby the reference of a demonstrative element is fixed by the *context* of the perception in which it occurs.³³⁹ One prospective counterargument to this is that the demonstrative element in question should not be taken to be involved in thought about the object demonstratively referred to (i.e. the cup), but as being some category of perceptual representation that is more basic than such thought, i.e. it forms non-conceptual content.³⁴⁰ I shall not rehash the concerns with non-conceptual content discussed in Chapter 1, though they would apply here too. A further concern relevant to this particular case, though, is that shifting the demonstrative element to the domain of non-conceptual content does not sufficiently heighten the explanatory role it can play. This argument has made the demonstrative element itself available to the subject by building it in to their subjective life at the level of non-conceptual content. What fixes the reference of this demonstrative element are certain aspects of the context of the perception in which it occurs, and these are not made available to the subject on this account. Consequently, how the subject can come to have any understanding of what the demonstrative element refers to – or even that it refers at all – is arguably left wholly opaque on this picture.³⁴¹

It has been taken to be the case that representationalism cannot account for our capacity to refer demonstratively, largely on the basis that the latter involves a close connection to the mind-independent object referred to. For the sake of covering all bases, I shall close this line of argument by briefly assessing the prospects for the representationalist resisting this conclusion by denying that this close connection obtains. We saw in the preceding paragraphs problematic representationalist efforts to accommodate our capacity for demonstrative reference that jeopardised the connection between the subject’s experience and the object demonstratively referred to.

³³⁹ See Burge (1991)

³⁴⁰ See Campbell (2002: 125 – 126)

³⁴¹ Ibid.

The problem with denying that demonstrative reference involves a close connection to a particular mind-independent object at all is that it leaves the demonstrative reference as a whole without anything onto which to anchor. In summary: this approach instead jeopardises the very connection that we take demonstrative reference to involve – it is wholly unclear, on this final approach, what (if anything) is being demonstrated to at all. There is nothing that serves as an anchor that grounds what ‘*that* object’ actually picks out.

I have considered in this section the notion of demonstrative reference, particularly with respect to the es-question (REL-2):

(REL-2): What makes demonstrative reference possible?

I submitted above that if it could be shown that (i) demonstrative reference inherently involved a close connection between the subject doing the referring and a mind-independent object, and (ii) that relationalism could accommodate this connection whilst representationalism could not, that (REL-2) would constitute an es-question to which relationalism but not representationalism could provide an eg-answer. I take the foregoing to have satisfied (i) and (ii), as relationalism can provide an eg-answer via the perceptual relation of acquaintance, whilst representationalism seemingly cannot accommodate it at all, and denying (i) apparently comes at the price of losing the grounding for our capacity for demonstrative reference. I conclude, therefore, that (REL-2) is such an es-question.

4.2: *The interests of representationalism*

I turn now to identifying an es-question that is eg-answerable by representationalism but not by relationalism. In the following section, I shall accomplish this by considering non-veridical experience; however, since representationalism is primarily a theory of veridical experience, I shall attempt now to find such an es-question by considering veridical experience. Drawing on Dan Cavedon-Taylor’s (2015), which considers perception of

properties relating to what *kind* of object that object is classified as (i.e. kind properties), I submit that a good candidate es-question is:

(REP-1): How do subjects of veridical experiences perceive the kind properties of mind-independent objects?

In what remains of this section, I shall provide an outline of the ideas and arguments Cavedon-Taylor presents for thinking that addressing the question of how we perceive kind properties weighs in on the side of representationalism, which I take to constitute an argument against the view that relationalists can offer a satisfactory eg-answer to the es-question that is now under consideration.

First, we start with the basic claim

(KP): subjects can perceptually experience the kind properties of objects.³⁴²

Though perhaps not itself uncontroversial, if we take (KP) for granted, and we take representationalism and relationalism to be the two front-running theories in the philosophy of perception, *and* deny their compatibility, then a natural question that emerges is to determine which of the two theories best accommodates (KP). From here, we have two different adaptations of (KP), that differ because of the distinct commitments of relationalism and representationalism:

- 1) **(^{Rep}KP):** *Subjects can visually represent kind properties to be instantiated by objects.*³⁴³
- 2) **(^{Rel}KP):** *Subjects can visually detect the kind properties of objects.*³⁴⁴

For relationalism, then, it appears the properties of natural and artefact kinds (to unpack the umbrella term of 'kind properties', more on this below) will need to be such that they are candidates for visual detection, where to detect something visually is to become perceptually acquainted with it. However, such properties are arguably exceptionally *poor* candidates for

³⁴² Cavedon-Taylor (2015: 487)

³⁴³ Ibid. (492)

³⁴⁴ Ibid. (494)

visual detection. This is because bare visual detection is a direct detection, with limited contribution from the subject themselves. The information picked up through visual detection is information about the subject's environment that does not rely heavily on subjective mental interpretation governed by convention or similar; it is, as it were, information that is wholly 'out there in the world'. Natural kind properties (e.g. 'being water'), meanwhile, are 'hidden essences', of the sort that render the object in question to be *thought of* or *classified as* the object it is; artefactual kind properties (e.g. 'being a guitar'), conversely, are responsive to their creator's intended function for them to serve.

In each case, it seems, the subject's interpretive capacities are utilised: one must learn to appreciate the various criteria for 'being water', though one could visually detect 'that stuff over there', for example.³⁴⁵ Though there may be some theory of natural or artefactual kinds according to which it is viable that these sorts of properties can be visually detected (in the sense above), it is argued that not only would committing to such theories of kinds constitute a high theoretical cost, but also that it remains unclear as to what such a theory would look like.³⁴⁶ By way of example, here is Cavendon-Taylor:

*In looking at a doorknob, the relevant properties that make it a doorknob are such things as: the property of having been made by a creator with such and such intentions, the property of having such-and-such a function, the property of having been made in such-and-such a social/cultural context, etc. These are not properties to which one can be perceptually related. After all, these properties are, themselves, relations to a particular person, a function, and a context.*³⁴⁷

³⁴⁵ Natural kinds, and in particular the example of water (as H₂O) has been hotly contested at least since Putnam (1975). The literature on the topic is massive, and I cannot hope to engage with this separate issue here, due to constraints on the scope of the present project. Bird & Tobin (2017) provide an excellent introduction to the debate and different views people have on it. With that caveat in place, I commit myself without defence in what follows to a certain position on kind properties with which not all parties to the debate would agree. Two thoughts: if one objects particularly strongly to this, one can disregard (REP-1), and I think my argument as a whole still works; as we shall see, however, the representationalist account of (KP) to be considered is compatible with a number of positions on kind properties – *including anti-realism* – so this position is not as weak as it may appear, and I insert the above caveat only for the purposes of transparency.

³⁴⁶ Ibid. (497)

³⁴⁷ Ibid. (498)

In this example then, the ‘relevant properties’ are just those properties that make a doorknob a doorknob – it is just these properties that constitute the kind properties of a doorknob, indeed of *any* doorknob: the properties enjoyed by all and only doorknobs. In the case in question, we have considerations such as specific function for which doorknobs are made, as well as intention with and context in which they are manufactured. As Cavedon-Taylor points out, these are not properties to which one can be perceptually related in the direct manner involved in visual detection, as they are themselves relations to particular functions, people and contexts that one does not detect simply by visually detecting the object itself.³⁴⁸ It seems, then, that kind properties are not good candidates for visual detection and – if relationalism holds that visual perceptual acquaintance is a matter visual detection – then (^{Rel}KP) seems unattractive.

Of course, this falls short of assuring us that (^{Rep}KP), and with it representationalism, fares any better. On the face of it, representationalism seems better equipped. This is because the view separates *seeing* the objects in one’s environment on the one hand, and the *attribution* of properties to those objects on the other (i.e. autorepresenting) – the former is a process which operates from the outside of the subject inwards; the latter operates from within the subject outwards.³⁴⁹

In this respect, seeing the objects is – as it were – a ‘bottom up’ process, where information about the physical environment culminates within the subject’s conscious experience of their environment. Conversely, the attribution of properties would be a ‘top down’ process, where various properties conceived of by the subject are projected outwards and assigned to objects in their physical environment. Call this position *separatism* about seeing and attributing properties. In endorsing separatism, then, the representationalist successfully severs questions regarding *what* properties

³⁴⁸ Ibid. (498)

³⁴⁹ See earlier discussion (§1.3.4) on autorepresentation and allorepresentation for arguments in support of my view that representation works this way.

the perceiver experiences from questions about the *nature* and metaphysics of those properties.³⁵⁰

This latter point is quite telling, as it implies that *even if* there could be some theory of kind properties such that they were candidates for visual detection, thereby vindicating (^{Rel}KP) in the manner described above, to embrace (^{Rep}KP) is still to tread the path of least resistance. This is due to representationalism's embrace of separatism: no theory of kinds must be assumed in order to accommodate the thought that we can perceive kind properties – the representationalist interpretation of (KP) is not a matter of being related to kind properties (and therefore of having to provide a robust account of their metaphysics), but rather is a matter of perceptually representing those kind property instances or, in other words, of *attributing* them to mind-independent objects.³⁵¹ This thought can be drawn out with reference to the familiar natural kinds example of water being H₂O. Put simply, it is flatly a mistake to assume that possession of the concept H₂O is required for the representationalist to see water *as water*. Rather, on a representationalist account, it is not required in order for that to happen that the subject represent water in a distinctly scientific way such as this: “[p]ossession, and subsequent activation of, the concept WATER can do the job instead.”³⁵² The key point being made here is that *whatever* way the subject understands the kind properties of water, they can autorepresent some substance in their environment as instantiating that property, and do so without thereby being committed to providing a robust account of the metaphysics of these kind properties.

As a final point, note too that adopting an anti-realist position about the kind properties in question *still* favours the representationalist account of (KP) over the relationalist one. This is because if anti-realism about kind properties is true – and objects really don't have those 'hidden essences' that make them the objects they are – then properties such as 'being water' or

³⁵⁰ Ibid.

³⁵¹ Ibid.

³⁵² Ibid.

'being a tomato' are incapable of instantiation; however, whilst it follows from this that they cannot be visually detected (and so (^{Rel}KP) comes out false) it is still possible for the subject to perceptually attribute kind properties to mind-independent objects *even though* there is *de facto* (on the anti-realist account) no way in which they could ever actually be instantiated by those objects.³⁵³ Whilst this admittedly would amount to an erroneous perceptual experience, this is arguably not of as much consequence as it may seem. Cavedon-Taylor again:

*[A]n error-ridden account of KP is no more incoherent than is, say, an error-ridden view of colour perception, no matter its unattractiveness. By contrast, in understanding perceptual experience to be a relation, the naïve realist cannot recognise the possibility of an error-ridden theory of KP; for them, KP entails commitment to the existence of kind properties themselves, whereas for representationalism, KP does not.*³⁵⁴

It is on these bases that the thesis (KP) – that we can perceptually experience kind properties of objects – is taken to favour representationalism over and above relationalism. Indeed, representationalism looks set to give a better account of our experiences in this regard even if an anti-realist approach is taken towards kind properties in general. Consequently, I take it that the representationalist, though not the relationalist, has at their disposal a satisfactory eg-answer to (REP-1).

Admittedly, the case I have made falls short of suggesting that the relationalist *cannot possibly* provide such an answer, or an objection to the foregoing analysis. An instance of the latter that I anticipate is that the relationalist may appeal to conceptual registration, as discussed earlier in this chapter, and in Chapter 2. The idea being that kind properties involve concepts and registration and deployment of these actually *is* a central part of the relationalist picture. Against this, I would reply that I have already argued (see §4.1.1) that conceptual registration and the role it plays for relationalism strikes me as very similar to representation. If this is so, then I

³⁵³ Ibid. (499)

³⁵⁴ Ibid.

think the case is there to be made that representationalist apparatus apparently plays a fundamental, indispensable role in explaining our perceptual experiences. Whilst I am happy to concede this point, the austere relationalist should not be. I think this is the case more generally, but if relationalists wish to appeal to conceptual registration to accommodate our perceptual experiences of kind properties, then I submit that this merely strengthens the point.

Additionally, if the above analysis of the situation is adequately correct and comprehensive, then it seems as though efforts to provide an alternative relationalist eg-answer to the one I have envisaged would still come at the cost of biting a substantial bullet in the form of a theory according to which kind properties are candidates for visual detection – a claim which would then require some undesirably strict constraints as to the metaphysics of the kind properties in question.

4.3: Further examples: non-veridical experiences

Finally, I wish to discuss both types of non-veridical experiences – illusion and hallucination – in order to identify further examples of es-questions answerable only by one or the other (but not both) of representationalism or relationalism. I shall revisit a point touched on in Chapter 2, pertaining to the thresholds for error according to which a given non-veridical experience is categorised as either an illusion and hallucination. The point highlighted that illusion definitionally involves a misperceived object, but an object nonetheless, and maintained that this placed constraints on how erroneous an illusory experience could be. By contrast, hallucinations do not do this, as there is no mind-independent object appropriately causing the non-veridical experience. I think this point plausibly demonstrates an instance of the relationalist view alone offering a satisfactory eg-answer to a legitimate es-question.

I take it to be fairly uncontroversial on definitional grounds that unlike hallucinations, which might be of anything whatsoever, illusions *must* be restricted in some way in how erroneous they can be, due to their object-involving nature. Furthermore, because relationalism takes seeing to be basic – in the sense that successful veridical perception is the foundation of the account of experience offered – the view allows a derived account of mere appearances to be constructed. If seeing that an object *o* looks F is a matter of standing in a relation to *o* from a given standpoint *c*, which are such that – to the subject – *o* appears F on account of standing in this three-place relation, then the illusory case can be constructed in a similar way, except that *o** is such that it appears F despite being G, not F, and this too is due to the orientation of *o* relative to the subject, and to how *c* as a third relatum influences the acquaintance relation that holds between *o* and the subject.

In other words, that *o** appears F when it is in fact not F but G is accounted for because *c* is such that objects like *o**, seen from the subject's perspective within *c*, look G. One can imagine *c* involving lighting conditions that cause misleading appearances and erroneous colour judgements of cars at night under streetlights, for instance, just as one can imagine that spatial orientation of the subject and object are such that the subject's view of something salient to the error is occluded. The point is that the error involved in illusory experience, whatever it is, is explicable by these facts of the situation; the subject's point of view, or some other fact about objects like *o** when viewed in *c* accounts for *o** appearing F when it is in fact G.

Since these facts about the three-place relation involved in an illusory experience are taken to wholly explain the illusion, the nature of the error committed by the subject of an illusion is wholly answerable to these facts. By contrast, a hallucination cannot be explained in this fashion: since the experience is not appropriately caused by a mind-independent object, there is no three-place relation to serve as the foundation for that sort of account of hallucinatory experience. Consequently, the nature of the error committed by the hallucinating subject is not wholly answerable to facts about how the world is, contrasting with the illusory case. Both of these observations agree

with our intuitions about the types of non-veridical experience, based on the rudimentary definitions provided of each.

Another way of putting this insight is that illusion's world-involving nature places a restriction on just *how* erroneous the experience can be, relative to the fact of the matter, before it is rightly classed as a hallucination.³⁵⁵ If the foregoing analysis is correct, then relationalism can account for this, and thus can provide an answer to the following:

(REL-3): How do we determine whether an erroneous experience is illusory or hallucinatory?

Conversely, because representationalism treats the way things *appear* or *look* to be as basic, the strategy of founding an eg-answer to this question by deriving an account of the mechanics of illusory experience from the mechanics of a more basic account of veridical perception is unavailable to its proponent. As a result, the thought goes, they cannot appeal to the workings of veridical experience to place the limitations on the error involved in a non-veridical experience that are necessary for determining whether it is illusory or hallucinatory.³⁵⁶

To be clear, I am not denying that representationalism can offer an account of how things appear to a subject to be – I take it that autorepresentation allows the subject themselves to represent things to themselves as being a particular way, and it is this particular way that is how things will appear to the subject of that experience, irrespective of the veridicality of the experience. What I am denying, because this representationalist account of appearances is not derived from a separate, more basic account of seeing, is that the account in question can provide any limitations on the erroneousness of illusory experiences.

Indeed, since representationalism cashes out both veridical perceptions and mere appearances in terms of content, it is not clear exactly to what

³⁵⁵ Brewer (2017: 217 – 218)

³⁵⁶ *Ibid.*

theoretical apparatus its proponent can appeal to appropriately restrict the degree of error involved in illusory experience of an object, except for experiential content³⁵⁷. Presumably, it will need to somehow involve representational content. One can imagine a response running along the lines of the experiential content of the illusory experience being *caused* by the mind-independent object(s) observed, contending that the content involved in this causal relation is the content of a misrepresentation.

The problem with this approach is that it commits the representationalist to explaining how this causal relation alone generates certain errors in the experiential content that are compatible with observing the object: in other words, the bare relation of causation must account completely for the misrepresentation. More plausible, perhaps, is an appeal to whatever mind-independent object has caused the erroneous experiential content, together with the relation this stands in to other worldly states of affairs. The issue with this approach is that it is essentially the relationalist story, with the added level of experiential content. This makes the inclusion of the latter arguably superfluous and unacceptably *ad hoc*.

More precisely, adding in the experiential content contributes nothing to the explanation: the heavy lifting is still done by the mind-independent objects themselves, their surface reflectance properties, lighting conditions in the environment, spatial orientation with respect to other objects, and so on. It may be the case that the experiential content of any illusion is, as a matter of fact, appropriately restricted in order to match closely enough with the (mis)perceived mind-independent object for the experience to be classed as illusory rather than hallucinatory. If that is what is to be argued, then the representationalist must provide an account of this that does not yield to the relationalist picture, otherwise the resulting eg-answer cannot be said to be representationalist.

³⁵⁷ I use the term experiential content as a catch-all term for all content centrally involved in perceptual experiences – representational or misrepresentational.

Representationalism takes content to be a feature of any kind of perceptual experience: veridical, illusory or hallucinatory. However, if the above is correct, then the view does not seem to do the job of outlining how erroneous an illusion can be, or of determining where the threshold – in this respect – exists between illusion and hallucination. If so, then it can be concluded that relationalism, but not representationalism, can provide an *eg*-answer to (REL-3).

However, even if relationalism is plausibly the only theory that can explain how some threshold for error might be put in place for illusions, I submit that representationalism is the only theory offering an explanation of potential subjective indistinguishability between veridical perceptions and hallucinations. If this is right, then it also shows how the two views can work together by explaining different aspects of non-veridical experiences that we would like to be explained. Relationalism alone may be able to explain precisely how erroneous a non-veridical experience can be before it is classed as a hallucination and not an illusion, whilst representationalism alone can provide an illuminating account of why veridical perceptions and subjectively indistinguishable hallucinations seem to share a common phenomenal character.

Switching from illusion to hallucination, we discover what I take to be the counterpart to the example of the previous paragraphs: representationalism provides an *eg*-answer to an *es*-question that relationalism does not. The *es*-question I have in mind is the following:

(REP-2): Why are certain hallucinations and veridical perceptions subjectively indistinguishable?

Rather than focusing on the representationalist's account of hallucination, though, it is easier to get this point across by sticking with the relationalist account for now. The standard relationalist approach to hallucination, as discussed in Chapter 2, is the so-called *epistemic* account of hallucination. According to this account, all there is to say in explaining hallucination at all is that one is in a state of perceptual experience such that one cannot know,

by reflection alone, that one is not undergoing the corresponding veridical experience. This wholly negative epistemic claim constitutes the most widely endorsed metaphysical disjunctivist analysis of hallucination, which is appealed to by relationalists because they require a way to resist the argument from hallucination, and can find this in a satisfactory account which holds that veridical perception and hallucination involve fundamental differences in mental state.

If the negative epistemic claim exhaustively accounts for the phenomenal character of hallucination, then all relevant es-questions that we might legitimately ask about it (including (REP-2)) must be provided satisfactory eg-answers by the claim itself, or that which underpins it. As we saw in Chapter 2, Martin's claim finds support in the form of the impersonal notion of indistinguishability to which the account appeals. Recall that for two experiences to be impersonally indistinguishable from one another was for them to be utterly incapable of being distinguished, rather than their indistinguishability being due to some failing or limitation on the subject.

Considering (REP-2) and what it asks, an eg-answer shall provide an account of why it is that a hallucination and its corresponding veridical perception are subjectively indistinguishable. A natural response to this question is that it has something to do with the phenomenal character of the hallucination, and its relation to the phenomenal character of the corresponding veridical perception. A satisfactory eg-answer would also be able to say what it is about phenomenal character that renders the two experiences subjectively indistinguishable.

The problem is that Martin's view cannot appeal to any robust positive feature of phenomenal character: neither of the hallucination, as he thinks there is none, nor of the corresponding veridical experience, as this would not really explain why the two experiences are subjectively indistinguishable. Martin has appealed to impersonal indistinguishability in order to account for the phenomenal character of hallucinations (see §2.4.1), yet doing this here would beg the question, for (REP-2) stipulates that the subjective

indistinguishability of the two experiences is exactly what the eg-answer should be explaining!

If this is correct, then Martin's metaphysical disjunctivist account seemingly cannot provide an adequate eg-answer to (REP-2). Representationalism, by contrast, will simply appeal to the common factor approach it advocates when considering hallucinations and corresponding veridical perceptions, arguing that both experiences share phenomenal character because both experiences share content, which is the result of successful representation in the good case and of misrepresentation in the bad case. In this way, (REP-2) is a legitimate es-question to which the popular relationalist account of hallucination – Martin's metaphysical disjunctivism – cannot provide a satisfactory eg-answer, whereas representationalism can.

There are potentially ways for the relationalist to avoid this unfortunate outcome for their view regarding (REP-2). If they can provide an alternative account of hallucination, then they will not by necessity be committed to the problematic outcome that Martin's disjunctivism is. Such alternatives would have to be shown to provide an eg-answer to (REP-2), however, and would need to be such that they weren't inherently problematic. I will now examine two such alternatives, discussed earlier in Chapter 2. The first is the alternative metaphysical disjunctivism advanced by Fish; the second is McDowellian epistemological disjunctivism.³⁵⁸

Recall that Fish's view was such that philosophers' hallucinations (those mundane hallucinations that are subjectively indistinguishable from counterpart veridical perceptions) lack phenomenal character, and explains their apparent possession of it in terms of false beliefs founded upon corresponding veridical perceptions, and the phenomenal character possessed by the latter experiences. It seems that this view, on the face of it, is in a position to provide an eg-answer to (REP-2). The eg-answer will essentially involve the cognitive effects of the hallucination in the manner outlined in the articulation of the view. Certain veridical perceptions and

³⁵⁸ See Fish (2009) and McDowell (1982), as discussed in Chapter 2 (§2.4.1 and §2.4.2 respectively).

hallucinations are subjectively indistinguishable because subjects of the hallucinations form the mistaken belief that they are having an experience with phenomenal character, and the phenomenal character they think their experience has is just the phenomenal character that the counterpart veridical perception would have.

We also saw some problems with the view, however. In the first instance, Fish's account is committed to upholding the seems/is distinction for phenomenal character. Generally, if what an experience is subjectively like for me is p , then p is the phenomenal character of that experience. Fish's eliminativism about phenomenal character for hallucinations flies in the face of this. More pressingly, we saw that once the philosophical conception of hallucination is understood as including philosophers' hallucinations, drug-induced hallucinations and hallucinations of impossible figures, the view leads to an intolerable disunity in accounting for different types of hallucination.

As a result of this, it is arguable that the view is not satisfactory due to inherent flaws. Consequently, even if it can provide an eg-answer to (REP-2), these flaws prevent it from serving as a satisfactory account of hallucination for independent reasons. Moreover, those independent reasons quickly signpost other potential es-questions to which the view could not provide a satisfactory eg-answer, such as 'how can we hallucinate impossible figures?'.³⁵⁹

The final way for the relationalist to provide an eg-answer to (REP-2) is to drop metaphysical disjunctivism as the account of hallucination, and to accept epistemological disjunctivism instead. Recall that both of these accounts involve the claim that one is in a different epistemic state when one is hallucinating than one is in when veridically perceiving.³⁵⁹ Metaphysical disjunctivism develops the epistemic claim in various ways, and then concludes that that is all there is to say about the phenomenology of hallucination. Epistemological disjunctivism does not make this commitment. In this respect, the door is open for some view or other to

³⁵⁹ See §2.4.2.

provide an eg-answer to (REP-2) on the basis of epistemological disjunctivism.

I am sympathetic to this approach, yet only because I think it allows a pleasing story to be told about hallucination that involves both relationalism *and* representationalism. Epistemological disjunctivism is an account of the *epistemic* aspect of hallucination. Whilst this leaves the phenomenology free to be explained by one's preferred view, epistemological disjunctivism *itself* does not have such ambitious goals. I surveyed in §2.4.1 and in the preceding paragraphs other wholly relationalist efforts to account for the phenomenology of hallucination, and have found them to be inadequate. Representationalism, conversely, has always been able to answer (REP-2) by appealing to representational content. Since epistemological disjunctivism makes no claim about the phenomenology of hallucination, it seems that the two are not inherently incompatible, and a hybrid theory of sorts might allow the relationalist to provide an eg-answer to (REP-2).

Of course, the view that falls out of this accepts that relationalism and representationalism are indeed compatible. An application of *modus tollens* informs us that austere relationalists are therefore unlikely to accept the view suggested. I have surveyed Fish's account of hallucination, and have argued that it cannot be relied upon due to inherent flaws which, if one does utilise it to generate an eg-answer to (REP-2), merely serve to indicate alternative es-questions that the view cannot answer. I have also evaluated Martin's metaphysical disjunctivism, which is also relied upon by other relationalists, such as Brewer.³⁶⁰ This view also was unable to provide an eg-answer to (REP-2), because on the one hand his view cannot by its own lights appeal to any positive feature of hallucinatory phenomenology, and on the other hand a satisfactory eg-answer to (REP-2) cannot appeal baldly to subjective indistinguishability, as that is precisely what we are trying to explain. If all of the above is correct, it can be concluded that

³⁶⁰ Brewer (2011: 109)

representationalism is the only theory of perception in the vicinity that can provide a satisfactory eg-answer to (REP-2).

If the rather large number of arguments and ideas considered in this chapter have fallen into place successfully, then I take it to have been shown that there exist at least several es-questions, such that with regard to representationalism and relationalism, each can be offered a satisfactory eg-answer by only one of these theories, but not by both. On the side favouring relationalism, we have:

(REL-1): How does experience provide the perceiving subject with *conclusive* reasons for applying certain empirical concepts in judging their environment as being a particular way?,

(REL-2): What makes demonstrative reference possible?, and

(REL-3): How do we determine whether an erroneous experience is illusory or hallucinatory?

On the side favouring representationalism, meanwhile, we have:

(REP-1): How do subjects of veridical experiences perceive the kind properties of mind-independent objects?, and

(REP-2): Why are certain hallucinations and veridical perceptions subjectively indistinguishable?

In the broader dialectic, this would mean that relationalism and representationalism, *qua* overall explanations of perceptual experience, are such that each explanation offers something that the other does not. As I shall show in Chapter 5, this will be extremely important in arguing that relationalism and representationalism are *complementary* explanations of one and the same phenomenon: perception.

5: *On Complementary Explanations*

In Chapter 3, I argued that we should understand explanations as consisting in relations of explanation-seeking questions (*es-questions*) and explanation-giving answers (*eg-answers*). I have supported the methodology of *explanatory pluralism*, according to which there are no rules of preference or exclusion for explanations that are based on the types of explanation they are. Rather, once we identify an explanation as an *eg-answer* to some *es-question*, we find that the *es-question* (possibly with some refinement) indicates the explanatory interests of the inquirer. These, in turn, provide the necessary context within which to evaluate explanations for which is to be preferred and which is to be excluded. The *eg-answer* which proves most illuminating relative to the context determined by these explanatory interests simply is the one to be preferred, according to explanatory pluralism. I presented cases where such an approach is motivated in other disciplines, and then considered the prospective merits and motivations for applying it in the philosophy of perception.

Then, in the immediately preceding chapter, I utilised this explanatory pluralist methodology and *es-question/eg-answer* framework for explanations in the relationalism versus representationalism debate to identify certain legitimate *es-questions* – ones to which we would plausibly want *eg-answers* to be given by our philosophical theory of perception. These *es-questions* were such that of the two theories of perception under discussion, only one of them could provide an *eg-answer*. I argued that multiple *es-questions* of this nature could be found that favoured each theory, claiming that these *es-questions* would be of paramount importance to my efforts to reconcile relationalism and representationalism.

Now that I have identified these *es-questions*, I am in a position to construct my argument against the monist position that relationalism and representationalism are incompatible. This will be done by observing ways in

which allegedly competing explanations might interact. In particular, I will discuss the notion of explanations being *complementary* to one another in various ways. Following Marchionni, I will spend the first half of this chapter elucidating precisely what it means and what it takes for explanations to be either *strongly* complementary or *weakly* complementary. Thereafter, I will utilise the es-questions identified in Chapter 4 to support my argument that relationalism and representationalism *are* complementary explanations. In particular, it will be argued that these es-questions support a view which holds that they are *strongly* complementary explanations – in that their integration results in a better explanation overall – but in which the notion of *weakly* complementary explanations (to be defined in the following section) still plays a vital role.

5.1: *Strongly and weakly complementary explanations*

The notion of complementary explanations to be used in what follows would benefit from being elucidated precisely at the outset. Following Marchionni, I distinguish between two sorts of ways in which a pair of explanations can be complementary:

Weakly complementary: Two explanations of the same phenomenon are (i) legitimate explanations of that phenomenon, (ii) autonomous, and (iii) are complementary in virtue of possessing different explanatory virtues.

Strongly complementary: Two explanations of the same phenomenon are (a) legitimate explanations of that phenomenon, (b) each possess different explanatory virtues, and (c) when integrated, they provide a better explanation (in the sense of making more complete or enhancing the explanatory virtues of the individual explanations).³⁶¹

There are three claims being made in the definition of each sort of complementary explanation, which I shall now briefly elaborate upon. The first condition on being either sort of complementary explanation is that

³⁶¹ Marchionni (2008: 315); I have adapted the definitions offered in the original for the sake of clarity on what is to follow.

each of the explanations of the same phenomenon be a *legitimate* explanation of that phenomenon. By this, I simply mean that each offers a satisfactory eg-answer to an es-question that falls within the explanatory interests of those wishing to understand that phenomenon.

I now address the remaining conditions in turn. For *weakly* complementary explanations, the second condition is that the pair of explanations are *autonomous*. What is meant by this is simply that the explanations are not integrated, and proceed as standalone explanations. This could be because there is nothing to be gained from integrating them into one single explanation. A pair of explanations could remain autonomous because they each offer a different eg-answer to the same es-question. In this respect, the two explanations would offer competing answers. According to the explanatory pluralist methodology I am employing, which eg-answer is to be preferred will depend upon the inquirer's explanatory interests. Importantly, as I shall argue later, a pair of explanations offering competing eg-answers like this does not entail that the explanations as a whole are in competition, and so does not preclude their being complementary, or even compatible.

The third condition requires more unpacking. Before doing so, it is worth clarifying some terminology. What I will be considering here are explanations of the same phenomenon which are complementary. The single phenomenon is perception. The explanations in question are relationalism and representationalism. Each of these explanations consists of a set of eg-answers to a multitude of fine-grained es-questions about perception. Taking this set of eg-answers together results in an eg-answer to a more coarse-grained es-question about perception, such as 'what is the nature of perception?' or 'what philosophical account of perception should we give?'

I return to the third condition on *weakly* complementary explanations. In the first instance, it states that weakly complementary explanations possess different explanatory virtues. I take an explanatory virtue to be a property of an explanation or theory: if the explanation or theory is admirable in view of (a) its explaining something that another view cannot, or in view of (b)

explaining something in a particular way, then it possesses an explanatory virtue in this respect. The precise nature of explanatory virtues involved in satisfying (a) will depend on the particular explanations involved. Examples of the explanatory virtues involved in (b) are traits of explanations as a whole, such as *breadth* or *depth*. Breadth is concerned with unifying and accounting for several *explananda* via their commonalities; for instance, if representationalism offers an eg-answer that encompasses several es-questions (in the way that it offers representation as accounting for phenomenological aspects of both veridical *and* non-veridical experiences, say), it may be said to possess greater breadth, relative to this explanatory interest, than relationalism. Depth, conversely, is concerned with understanding a single *explanandum* by appeal to what sets it apart or distinguishes it from others; in this respect, an explanation may involve greater depth if the eg-answer it offers to a single es-question is more comprehensive and really gets to the heart of the matter.³⁶² An instance of this could arguably be relationalism offering an account of perceptual epistemology of veridical perception that gets to the heart of the matter by centrally involving mind-independent objects, in a way that representation does not. These are just examples; what is important to note presently is that the first of these involves comparison, whilst the second involves contrast.³⁶³

Moreover, the thought is that these explanatory virtues (in line with the autonomy condition) are – if not mutually exclusive – such that they trade off against one another. As I understand this, what is meant is that if we had a pair of explanations e_1 and e_2 , and suppose that they possess different explanatory virtues (say breadth and depth), then they are such that either (1) each one of them possesses all and only one of these explanatory virtues (i.e. the explanatory virtues are mutually exclusive with respect to these explanations), or (2) they each possess a mixture of breadth and depth, but are such that opting for e_1 will, say, involve a greater amount of explanatory

³⁶² Ibid. (319 – 320)

³⁶³ Ibid. (317)

breadth than depth, whilst opting for e_2 will yield a greater amount of explanatory depth *and* a lesser amount of explanatory breadth than is yielded by opting for e_1 (or vice versa), *relative to explanatory interests*.

It is worth noting that breadth and depth, as I define them above, are merely examples of explanatory virtues. What would now be helpful is to set the notion of explanatory virtues into the context in which I will be discussing them in what follows. In terms of our preceding es-question/eg-answer framework, establishing that two explanations possess different explanatory virtues is, I argue, achieved by demonstrating that each offers a satisfactory eg-answer to *distinct* es-questions.

The reason for placing importance on the es-questions being distinct, rather than simply the eg-answers, is that we are considering explanations of the same phenomenon – perception – which are complementary. The explanations in question are relationalism and representationalism, and each of these consists of eg-answers to a multitude of es-questions. If explanations simply produced different eg-answers to the same es-questions across the board, then the argument could be made that they are simply competing explanations. Note that this is not the same as maintaining that relationalism and representationalism are autonomous. Autonomy merely requires that they are not integrated, *possibly* because they offer competing eg-answers to *some* es-question(s). Maintaining that they are wholly competing explanations, on my account, requires that at least one of them *must* offer competing eg-answers to *all* es-questions that the other attempts to answer. I shall discuss this in greater detail below.

The remaining definitions pertain to *strongly* complementary explanations. The first of these states that the pair of explanations once more must possess different explanatory virtues, just as weakly complementary explanations do. The final condition is what separates strongly from weakly complementary explanations, and stipulates firstly that the explanations *can* be integrated (*contra* weakly complementary explanations), and that this integration leads to an enhancement in explanatory virtue (again, *contra*

weakly complementary explanations, whereby the explanatory virtues trade off against each other). One important example that is identified in Marchionni's work is that the integrated explanation could be more *complete* as a result, but in general the thought is simply that the explanatory virtues possessed by one of the explanations do not come at the expense of those possessed by the other, and in some respect the overall integrated explanation is possessing of some explanatory virtue that the unintegrated explanations do not.

Marchionni applies this strongly/weakly complementary distinction to explanations in the social sciences, yet there is nothing within the distinction as defined above that prevents it from being adapted and applied to explanations in other fields, such as the philosophy of perception. It is precisely this that constitutes my goal for the remainder of this chapter. First, however, there are a couple of other distinctions and pieces of terminology that Marchionni draws upon in her analysis of complementary explanations, and these must now be introduced.

The first of these is the distinction between micro- and macro-explanations. Though this distinction is, strictly speaking, one which will only obtain properly in explanatory enterprises similar to the social sciences, there nevertheless remains good cause for explicating the distinction here, as more than one element of its machinery will prove useful in what follows. To be precise: the term 'micro-explanations' is reserved solely as a way of characterising explanations which operate "exclusively in terms of individuals and their intentions"; conversely, the term 'macro-explanations' refers to explanations which operate "exclusively in terms of macro-aggregates or social structures."³⁶⁴ Explanations that are neither micro- nor macro-explanations, states Marchionni, are called "mixed" explanations.³⁶⁵

³⁶⁴ Ibid. (316)

³⁶⁵ Ibid.

With this distinction in place, it is possible to introduce a ‘visualisation of possibilities’ for allegedly competing explanations:

- 1) Macro-explanations are better explanations [than micro-explanations; or *vice versa*].
- 2) The macro- and micro-explanation are *strongly complementary*: they are legitimate explanations that possess different explanatory virtues and are such that, when integrated, they produce a better explanation.
- 3) The macro- and micro-explanations are *weakly complementary*: they are legitimate explanations that possess different explanatory virtues, but are autonomous to the extent that each is a standalone explanation, the explanatory virtues of which are either mutually exclusive or trade off against one another.

With this range of possibilities in mind, it is then concluded – much in the spirit of explanatory pluralism – that “which [complementary] explanation is preferred in a given context depends on what first motivates the explanatory enquiry.”³⁶⁶

In order to illustrate how this works, Marchionni utilises an example – due to Jackson and Pettit – on explanations of rising crime rate statistics:

The phenomenon to be explained is an increase in the crime rate. The macro-explanation explains it by citing an increase in the level of unemployment. The micro-explanation describes the precise changes of opportunities and motivations of specific individuals that led to a rise in the crime rate. The macro-explanation focuses on the similarities between the actual world and other possible worlds in which the crime rate increased, and tells us that under a wide range of individual-level circumstances and psychological profiles, the rise in unemployment ensures that the crime rate increases. The micro-explanation instead describes the actual causal process, and thereby focuses on the differences between the way in which the crime rate actually increased and the other possible ways in which it could have increased. Switching from information to questions, the macro-explanation is an adequate explanation for a coarse-grained explanandum such as ‘why did the crime rate increase rather than decrease or remain the same?’, whereas the micro-explanation is an adequate explanation for a much more specific object of explanation, ‘why did the crime rate increase exactly as it did rather than in another slightly different way?’ Because the macro-explanation abstracts away

³⁶⁶ Ibid.

from the specific details of the particular increase in the crime rate, it is broader or more unified than the corresponding micro-explanation. To see this, consider the case in which we have two regions both experiencing a rise in the crime rate as a result of an increase in the level of unemployment (and further suppose that the unemployment-crime relation is a stable regularity). The unemployment explanation provides a unified explanation of the two occurrences, whereas the micro-accounts in terms of particular changes of opportunities and motivations are extremely different in the two cases.³⁶⁷

In the manner outlined in the preceding chapter, then, we see here that different explanations can be offered for what at a sufficiently coarse level of grain looks to be – and can be described as – the same es-question, yet on closer inspection and at a finer level of grain are in fact distinct es-questions. Both the macro- and micro-explanations from the above example might be treated as answers to the coarse-grained es-question ‘why did the crime rate increase?’. As the example goes on to illustrate, however, the eg-answers offered by the macro- and micro-explanations more specifically correspond respectively to the es-questions ‘why did the crime rate increase rather than decrease or stay the same?’ and ‘why did the crime rate increase as it did rather than in another slightly different way?’. A reason for this disparity can be uncovered by being attentive to how exactly one arrives at micro- and macro-explanations respectively:

Macro-explanations [. . .] provide ‘modally comparative information’: they focus ‘on similarities between the actual world and other possible worlds; it takes us to a distance at which we can discern constancies across the actual way things are and the way things might have been.’ Micro-explanations [. . .] provide ‘modally contrastive information’: they focus ‘on the differences between the actual world and other possible worlds’ it homes in [on] the particularities of the actual case.³⁶⁸

In other words, providing a modally comparative explanation of an event or phenomenon involves examining other possible worlds that feature the same event or phenomenon and seeking commonalities

³⁶⁷ Ibid. (318 – 319)

³⁶⁸ Ibid. (317); the notions of modally comparative information and modally contrastive information, which are applied to macro- and micro-explanations respectively, are discussed more fully in Jackson & Pettit (1992).

(which, in theory, may likely account for the occurrence) between these other possible worlds and the actual world. Modally contrastive explanations of an event or phenomenon, on the other hand, are formulated by examining other possible worlds where the event or phenomenon could have happened but didn't, and analysing the specific differences between these worlds and the actual world. With all of this laid out, we can see why it might be suggested that "modally contrastive information is given when the object of the explanation is described at a very fine grain, and modally comparative information when the object of explanation is coarser grained."³⁶⁹ After all, modally contrastive information is mapped on to micro-explanations, as these provide eg-answers in terms of individuals and their intentions, which are themselves modally contrastive; conversely, modally comparative information is mapped onto macro-explanations, which abstract to a sufficient distance for allowing the discerning of constancies holding across individual differences that may obtain at a finer-grained level.

A point of key importance here, though, is that neither the micro- nor the macro-explanation is intrinsically better than the other, and each has its own selling points. The macro-explanation, for instance, is *broad* or *unified*, as in theory it demonstrates connections and common patterns in a manner supporting - in the spirit of Kitcher (1981) - the claim that "abstracting from the details of causal mechanisms can increase explanatory power."³⁷⁰ The micro-explanation, meanwhile, offers a *deep* explanation that looks to expose "the inner workings of the relevant causal mechanism", a feature which would render the micro-explanation the favourable choice if an explanation were to be valued - in line with the view of explanation offered by Salmon (1998) - solely upon how well it opens up 'black boxes' and shows their inner workings.³⁷¹ To clarify these explanatory virtues slightly: the broad explanation is providing modally *comparative* information, in the

³⁶⁹ Ibid. (318)

³⁷⁰ Ibid. (319)

³⁷¹ Ibid. (319 – 320)

sense that it focuses on making comparisons between the actual world and other possible worlds in which the same phenomenon occurs; in contrast, the deep explanation is providing modally *contrastive* information, in the sense that it focuses on differences obtaining between how the phenomenon to be explained might have occurred, and how it actually occurred.

In a manner familiar from the preceding chapter on the workings of explanatory pluralism, however, neither of these types of explanation are such that there is a principled decision procedure for either favouring or excluding either one of them; instead, the pluralist stance is adopted and – with respect to the distinction between broad and deep explanations just outlined, it is contested that “which [explanation] we prefer in a given context depends on whether we are interested in breadth or depth.”³⁷²

Having outlined Marchionni’s conception of complementary explanations, the ways in which competing explanations might interact with one another, and how this fits together with how explanations can be evaluated and selected as preferred explanations, all that remains to be done in this section is to draw the reader’s attention to the fact that, in what follows, the account I offer is designed to remain for the most part as ecumenical as possible viz. how some individual explanations are evaluated as better than others in terms of how well they address es-questions rendered salient by the inquirer’s explanatory interests. Different explanatory interests will render different es-questions salient, and therefore would render as better explanations those that may not be considered as so relative to different explanatory interests. What I argue this means is that there will be various ways of answering es-questions about perception, that broadly adhere to the tenets of explanatory pluralism, and are features of the complementary explanations relationalism and representationalism.

The exact respect in which the latter wind up being complementary – i.e. strongly or weakly, and why – will depend on whether the explanations are to be integrated or not. I contend that they should be, and that doing so will

³⁷² Ibid. (320)

result in an overall account of perception that is more complete than the unintegrated explanations. In this respect, I will be concluding that relationalism and representationalism are *strongly* complementary, and shall defend this position below. Overall, the finding that relationalism and representationalism are strongly complementary entails that the two *are* compatible after all, as it overcomes the screening off objection used to uphold the opposite conclusion.

In the following section, I will utilise the apparatus identified in this section in order to clarify how an explanatory pluralist methodology can be applied to the relationalism versus representationalism debate. This will provide a solution to the apparent tension between the allegedly competing explanations which maintains that, rather than being competing explanations, the two are in fact *complementary* explanations.

5.2: *The Argument from Complementary Explanations*

Recall that the screening-off objections, according to which a theory can be screened off as explanatorily redundant by another have, as part of their integral workings, the following premise:

Explanatory Redundancy: If x explains everything that needs to be explained, then y is explanatorily redundant.

where the variables are assigned either relationalism or representationalism, as the structure of the argument is itself neutral on which theory is preferred. It is this premise that is needed to get proponents of either austere view who wish to utilise the Screening Off Argument from the premise that their theory explains everything that we want a philosophical theory of perception to explain, to the premise that the allegedly opposing view is screened off as explanatorily redundant. The thought is that a screening off objection could be used by either austere relationalists *or* austere representationalists. It is this premise that I wish to resist, and I shall do so by illustrating that there are two readings of it, only one of which

would get the austere theorist to the conclusion they seek (i.e. that the competing view is explanatorily redundant), and then arguing – via what I am calling the *argument from complementary explanations* – that this one reading is such that it renders the above premise, Explanatory Redundancy, false (or a true but unsatisfied conditional) irrespective of which way round relationalism and representationalism are assigned to its variables.

My first task, then, is to elucidate the two readings of Explanatory Redundancy in question, and to show why one of these readings is not sufficient for establishing the conclusion of the Screening Off Argument. The distinction between readings that I have in mind pertains to the interpretation of ‘explains *everything* that needs to be explained’; in particular, it hinges on what exactly ‘everything’ is taken to be quantifying over. If we accept the reasonably uncontroversial claim that there are multiple aspects of perception that a philosopher of perception has interests in explaining, then the following two readings, I suggest, are plausible:

Interest-dependent: ‘Everything’ is restricted to certain explanatory interests / aspect of perception (e.g. the phenomenal aspects of perception – the subjective phenomenology of experience).

Interest-independent: ‘Everything’ is unrestricted as far as any legitimate *explanandum* in the philosophy of perception is concerned (so on this reading for theory *x* to explain everything is for theory *x* to offer an *explanans* that covers *all* legitimate *explananda* in the philosophy of perception).

With this distinction between readings carved out, the problem for proponents of the austere relational and representational views, I suggest, is that the interest-dependent reading makes Explanatory Redundancy at best a false conditional. On the assumption that relationalism/representationalism merely explains everything that needs to be explained relative to *certain* (i.e. not all) explanatory interests of the philosophy of perception, e.g. explaining the epistemological aspect of perceptual experience, then it is simply false that the consequent of Explanatory Redundancy follows.³⁷³ This is because

³⁷³ It might be contended that the assumption in this sentence can be rejected, but no austere theorist would reject this assumption regarding their own view, for this would entail that the theory under consideration

whatever theory the austere theorist wishes to establish as being explanatorily redundant may provide the most suitable eg-answer to a distinct yet equally legitimate es-question – or even to most or all of the remaining es-questions – and at this point it is false that the theory is explanatorily redundant. Consequently, it looks as though the proponents of austere views, if they wish to utilise a screening off objection, need to substantiate the interest-independent reading of Explanatory Redundancy.

The problem with establishing the interest-independent reading, I will argue, is that the antecedent then becomes much more difficult to satisfy and, consequently, even though the conditional claim of the premise comes out true and the argument as a whole is valid on this reading (which I take it to be), it fails to be sound as the further required premise – to the effect that the antecedent of Explanatory Redundancy is satisfied – turns out to be false. In this section, I will argue that relationalism and representationalism are explanations which – although traditionally taken to be competing with one another – are in fact strongly complementary, and therefore Explanatory Redundancy is only a true conditional on the interest-independent reading in virtue of having a false antecedent, thus allowing the overall argument supporting the Screening Off Argument to be overcome due to lack of soundness.

In order to establish that relationalism and representationalism offer complementary explanations, I adapt the ‘visualisation of possibilities’ offered by Marchionni regarding ways in which allegedly competing explanations might interact:

- 1) Relationalism/representationalism is the better explanation.

offered an explanation that addressed *none* of the es-questions pertinent to a philosophical account of perception successfully, which would be tantamount to arguing that their own view was itself explanatorily redundant. Of course, they may reject this assumption regarding the other austere view, however this is not only exceptionally uncharitable but, as we will see, will commit them to the unpalatable position that what appear to be perfectly legitimate concerns in the philosophy of perception are not in need of explanation.

- 2) Relationalism and representationalism are strongly complementary explanations: both are legitimate explanations, which possess different explanatory virtues and, when integrated, they produce a better explanation.
- 3) Relationalism and representationalism are weakly complementary explanations: both are legitimate explanations, are autonomous, and are such that they possess different explanatory virtues which trade off against one another.
- 4) Neither relationalism nor representationalism offers a satisfactory explanation.

Here, (1) obtaining involves declaring one theory or the other as outright explanatorily superior, to the detriment of the other, in the manner argued for by proponents of the austere theories, be they relationalists or representationalists. Most efforts to substantiate this claim involve an appeal to Explanatory Redundancy, as articulated above, however it should be noted that (1) may also be established if either relationalism or representationalism turns out to be untenable due to internal problems with the theory. (2) and (3) respectively claim that relationalism and representationalism offer explanations that are strongly or weakly complementary; I shall argue that (2) is most applicable to relationalism and representationalism, though (3) will also play a role in my overall account. In particular, I will maintain that the explanations offered by relationalism and representationalism as complete theories turn out to be *strongly* complementary, and submit that this suggests a reconciliatory view in the philosophy of perception is to be preferred to its austere counterparts. There will, on my account, be *parts* of the overall explanations offered by relationalism and representationalism that turn out to be *weakly* complementary. As we saw above when discussing what it was for two explanations to be weakly complementary, one of the conditions of this is that they must be autonomous (i.e. not integrated). I have argued that, though equally true for the case of representationalism and relationalism, this is not as damning for the enterprise of finding a reconciliatory view of perception as one might think, as it is importantly different to the claim that they are wholly competing explanations. Though (4) has no analogue in the visualisation of possibilities for competing micro-

and macro-explanations considered by Marchionni, I add it in here for the sake of exhaustively considering potential interactions between allegedly competing explanations. I shall now turn to considering what conditions must be met in order for each of (1) – (4) to be true.

Starting with (1), we have two ways of establishing this:

1a) Relationalism or representationalism is subject to Explanatory Redundancy.

1b) Relationalism or representationalism is inherently untenable.

We saw above that satisfaction of (1a) only achieves its goal if Explanatory Redundancy is read in its interest-independent form. In order to satisfy the antecedent of this reading of Explanatory Redundancy, recall, relationalism or representationalism must offer an *explanans* that covers *all* legitimate *explananda* in the philosophy of perception. With the outcome of the previous chapter in mind, I argue that this does not seem plausible for either view. The reason for this is that the five es-questions I identified over the course of that chapter ((REL-1), (REL-2), (REL-3), (REP-1) and (REP-2)) were instances where one theory could provide a satisfactory eg-answer to them, whilst the other could not. In these cases, the opposing theory seems to *not* provide an *explanans* to a legitimate *explanandum* in the philosophy of perception at all. If this applies to each of these es-questions, then I do not think the austere forms of relationalism or representationalism can satisfy (1a).

It may be objected at this point that, for instance, the austere relationalist may hold that (REP-1) or (REP-2) does not latch on to legitimate *explananda*, because the *explananda* in question simply are not legitimate, and so my argument does not go through. I think this is misguided. We saw in Chapter 4 that the es-questions under consideration were perfectly intelligible, and involve phenomena that do seem to occur, and which we would like our theory of perception to explain. The issue is not with the legitimacy of the *explananda*.

More promisingly, the austere theorist can concede the legitimacy of the *explananda*, whilst holding that their view does provide the correct *explanans*, and that the latter is deflationary or error-theoretic in nature. Though a more promising strategy, part of my analysis in Chapter 4 of each of these es-questions involved examining and assessing putative eg-answers offered by the opposing theory. In each case where this occurred, we saw that the latter were not satisfactory eg-answers to the es-question being considered. For example, my analysis of (REL-1) (which was about experience providing conclusive reasons) considered potential representationalist responses, none of which were satisfactory; my analysis of (REP-2) (on subjective indistinguishability of hallucinations and corresponding veridical perceptions) considered several relationalist accounts of hallucination, all of which were found wanting. Whilst it may be true, therefore, that the opposing theories do provide an eg-answers to some of these es-questions, such eg-answers were found to be unsatisfactory across the board. Being satisfactory, I take it, is a plausible desideratum without which this form of response loses efficacy, and so I take this type of response to my above argument to be ineffective, and (1a) is not satisfied by either austere view.

I turn now to (1b), which holds in the event that relationalism or representationalism is viewed as inherently untenable. This may be concluded if either view is internally inconsistent, or if it can be argued that one view or the other entails commitment to intolerably problematic consequences. Note that (1b) – were it to apply to both relationalism *and* representationalism – results in (4), according to which neither view offers a satisfactory explanation. For this reason, I will take the following argument to apply to (1b) and to most ways in which (4) might be satisfied. I also take this to be a suitable point of departure from the austere theories: perhaps these possibilities do spell real trouble for austere relationalism and austere representationalism, but as neither of those is the view I am seeking to defend, I do not assume their respective positions in what follows.

Instead, the view I defend can – and does – appeal to those parts of the austere theories that are rightly considered to be their real strengths: the eg-

answers offered to the es-questions listed at the end of Chapter 4. A corollary to this, as we saw in identifying those es-questions, is that the view is not committed to appealing to weaker aspects of either theory. In short: a reconciliatory view can pick and choose which parts of relationalism or representationalism it wants to accept and – more importantly at this point – can choose which weaker parts it wants to avoid. Consequently, for my purposes, the satisfaction of (1b) or even (4) by the explanations offered by austere relationalism or austere representationalism doesn't matter. This outcome does not entail any problem whatsoever for a reconciliatory account. Indeed, areas where the austere theories look to be intolerably weak may prove just as effective in outlining the correct articulation of a reconciliatory view as areas in which the austere theories look particularly strong. If one can pick and choose parts from either theory, the challenge becomes to identify areas where one *must* appeal to one or the other. Here, intolerable weakness seems just as viable an indicator as overwhelming strength.

I take the above to have addressed the possibilities of (1a), (1b), and most ways of satisfying (4). There is one way in which (4) may be satisfied that has not been addressed, and that is if a theory that is wholly different from representationalism and relationalism (i.e. not identical to, or made up of parts of them) was proposed, and this theory highlighted some hitherto unrecognised *explanandum* in the philosophy of perception that was sufficiently important that any theory that failed to offer an *explanans* would be viewed as not providing a satisfactory explanation. Presumably such an outcome would be equally problematic to the sort of reconciliatory view I am attempting to argue for, however in the absence of this new view's articulation – and therefore of further specifics – I am unable to address the possibility further, save for concluding that I do not take this to be an accurate reflection of how things currently are in the philosophy of perception.

5.3: *Relationalism & representationalism as complementary explanations*

If the above is correct, then of the range of possibilities (1) – (4) for allegedly competing explanations, I have provided reasons for accepting that representationalism and relationalism do not satisfy (1) or (4), subject to being able to successfully show that they satisfy (2) or (3). If it can be shown that they do satisfy (2) or (3), then it is true that they cannot satisfy (1) or (4), but I take the reasons provided in the preceding section to strengthen my overall case.

(2) claims that relationalism and representationalism are strongly complementary explanations, whereas (3) claims that they are weakly complementary explanations. I will be arguing that relationalism and representationalism satisfy (2), and are *strongly* complementary (though weakly complementary explanations will also have a role to play in the account to be proposed in what follows). The difference between the two lies in whether the complementary explanations can be integrated in a way that provides us with a *better* explanation overall. In the case at hand, the sense in which the explanation integrating relationalism and representationalism will be taken to be better will end up being – based on Marchionni’s reading with regard to strongly complementary explanations – a more *complete* explanation: one which bolsters the explanatory virtues of both breadth and depth, as defined earlier, across the board once the component explanations are integrated.³⁷⁴ What strongly and weakly complementary explanations have in common is that each constitutes a legitimate explanation of the target phenomenon, that each is a distinct explanation, and that each possesses different explanatory virtues. In what follows, I will first show that these common conditions can be met – such that it can be concluded that representationalism and relationalism are complementary explanations.

³⁷⁴ It is worth pointing out that Marchionni herself actively resists the claim that more complete explanations are *synonymous* with better explanations in an exclusive definitional sense. That is to say, the possibility is left open that an explanation may be in some sense better than another despite *not* being more complete; they may simply be *deeper* explanations, for example (2008: 330). The point here is that, so far as strongly complementary explanations have integration as an essential facet, that they produce explanations which are more complete is potentially a more probable outcome than other ways in which the resulting explanation might be better (by being more stable, for example; see Marchionni (2008: 318) for discussion).

Then, I will look to establish what sort of complementary explanations they are.

The question immediately in front of us is the following: in order for it to be the case that relationalism and representationalism are complementary explanations, what set of circumstances must obtain? A plausible three-point list that answers this question, I suggest, states with respect to the two theories that:

- a) Each explanation is a legitimate explanation.
- b) Each explanation is distinct.
- c) Each explanation possesses different explanatory virtues.

This list is achieved simply by deconstructing the definitions of strongly and weakly complementary explanations into the three individual criteria that they have in common. Further elucidation may be achieved, then, by clarifying the conditions under which (a) - (c) would be satisfied.

In order to be a legitimate explanation, I take it that the explanation under consideration must provide a satisfactory eg-answer to a salient es-question. This is a fairly minimal definition, and can be distinguished from a more demanding one according to which an explanation is only legitimate if it offers eg-answers to many, most, or even all of the es-questions that one might reasonably ask of the target phenomenon. I think the more minimal definition functions as a better determination of an explanation's legitimacy. If all theory x offers is a single (correct) eg-answer to a single es-question that one may raise about a given phenomenon, then it is still plausibly a legitimate explanation of that phenomenon, on a reasonable reading of legitimate. That theory y provides correct eg-answers to every es-question relevant to that phenomenon may be a reason to far prefer theory y to theory x , but it does not count against x 's legitimacy as an explanation. Perhaps one is only interested in the es-question to which x provides an eg-answer. In such a case, x may rightly be preferred to y , as the latter may

offer vast amounts of superfluous information that obfuscates the explanation initially sought.

Whilst each of these claims seems true, it also seems entirely arbitrary to declare the explanation not legitimate on the basis that it is only legitimate at a certain finer level of grain. If such a move were permissible, it is hard to see why *all* of the explanations offered by both representationalism and relationalism cannot similarly be deemed not legitimate explanations on the grounds that they do not, in virtue of the fact that the theories offering them are only attempting to offer accounts regarding our *perceptual* experience, provide eg-answers to all of the es-questions pertaining to, say, the nature of the mind as a whole.

The reading I intend for legitimacy to have here, regarding theories of perception, is that relationalism and representationalism are legitimate explanations of perception, and are so *with respect to the es-questions to which they provide eg-answers*. Since they both provide satisfactory eg-answers to es-questions that can reasonably be asked about the phenomenon of perception, they are legitimate explanations on my intended interpretation of 'legitimate explanation'. Having outlined what it is for (a) to be satisfied, in such a manner that I take it to have also been made clear that (a) in fact *is* satisfied, I turn now to (b) and (c).

In order for it to be the case that one explanation is distinct from another, I take it that the set of eg-answers offered by each explanation must not be identical. I place the importance on the side of eg-answers rather than es-questions here because this appears to be a more fruitful way of individuating explanations. Whilst it is true that two explanations that address distinct es-questions will be independent explanations, it is also true that two explanations that address exactly the same es-questions may be independent explanations, and this will happen precisely in the case whereby the same es-questions are addressed yet different eg-answers are provided. Cases of disagreement among competing explanations require that they are attempting to explain the same thing, yet offering different

explanations. In this respect, such disagreement is predicated upon them being distinct explanations. What cannot happen, since any eg-answer is an eg-answer to an es-question, is that two explanations are distinct despite offering all and only the same eg-answers as one another. No-one involved in the relationalism versus representationalism debate holds that they are not distinct explanations. Any number of differences between the views that I have discussed serves to satisfy this condition.

Finally, (c) holds that the two explanations must possess different explanatory virtues. Recall that explanatory virtues were defined as properties of an explanation or theory. This allowed for two sorts of explanatory virtues: those possessed by a theory because it is admirable in view of explaining something that another cannot, and those possessed by a theory in view of it explaining something in a particular way. Of these, the former will be essential to my argument in showing that relationalism and representationalism are complementary; once this has been established, both sorts of explanatory virtue will come into play when assessing whether they are strongly or weakly complementary explanations.

I take it that one way of demonstrating that explanations possess the first sort of explanatory virtue is to show that each explanation offers an eg-answer to at least one es-question *not* addressed by the other. For present purposes, this would amount to there being at least one pair of es-questions within the purview of the philosophy of perception, which are such that one of them is offered an eg-answer by relationalism but not by representationalism, whilst the other is offered an eg-answer by representationalism but not by relationalism. If this set of circumstances obtains, representationalism and relationalism possess different explanatory virtues (of the first sort described above).

Additionally, it seems as though if an explanation provides an *explanans* to an *explanandum*, then this plausibly is an explanatory virtue *relative to what we want to understand*. If there is some aspect of perception that we wish to understand and only relationalism (say) adequately explains it, then

relationalism possesses an explanatory virtue relative in view of offering this *explanans*. This is an explanatory virtue *in virtue of* explaining what it was that we wanted to have explained, and if two explanations can be said to possess explanatory virtues in this fashion, then they will possess different explanatory virtues (again, of the first sort).

This set of circumstances *does* obtain, and demonstrating that it does was the task with which the previous chapter was concerned. If one es-question in favour of either theory is all that is required to show that relationalism and representationalism possess different explanatory virtues, then the findings of Chapter 4 should more than adequately support the claim that they do. There, I identified three es-questions that favour relationalism in the manner required, and two that favour representationalism. For convenience, I re-state them here:

(REL-1): How does experience provide the perceiving subject with *conclusive* reasons for applying certain empirical concepts in judging their environment as being a particular way?

(REL-2): What makes demonstrative reference possible?

(REL-3): How do we determine whether an erroneous experience is illusory or hallucinatory?

(REP-1): How do subjects of veridical experiences perceive the kind properties of mind-independent objects?

(REP-2): Why are certain hallucinations and veridical perceptions subjectively indistinguishable?

In Chapter 4, I provided arguments for why each of these es-questions was such that it satisfied the condition of being eg-answerable by either relationalism or representationalism but not by both.

Appealing to the first sort of explanatory virtue defined above, I contend that relationalism therefore has explanatory virtue in view of explaining: (i) how experience provides subjects with conclusive reasons for applying certain empirical concepts in judging their environment as being a particular

way; (ii) how demonstrative reference is possible; and (iii) how we determine whether an erroneous experience is veridical or hallucinatory. We have seen that representationalism cannot provide satisfactory eg-answers to (REL-1) – (REL-3), and so relationalism – uniquely of the two views – possesses explanatory virtues in view of explaining these matters.

Applying the same rationale to the es-questions that favour representationalism, I contend that representationalism can be shown to have explanatory virtue in view of explaining (i) how we perceive the kind properties of mind-independent objects, and (ii) why certain hallucinations and veridical perceptions are subjectively indistinguishable. Again, Chapter 4 demonstrated that relationalism could not provide satisfactory eg-answers to (REP-1) and (REP-2), and so representationalism – uniquely of the two views – possesses explanatory virtue in view of explaining these matters.

If the above analysis is correct – if conditions on complementary explanations are accurate, the definition of explanatory virtue is acceptable, and my arguments to the effect that representationalism and relationalism possess different explanatory virtues whilst satisfying the other conditions are correct – then *relationalism and representationalism are complementary explanations*. The pertinent question then becomes whether they are strongly or weakly complementary. The answer hangs on whether they can be integrated in such a way as to provide a better explanation overall (in the sense of enhancing net explanatory virtue, i.e. being a more complete explanation), or are to be treated as autonomous explanations, with explanatory virtues that trade off against each other.

For representationalism and relationalism, that the first sort of explanatory virtues – the provision of eg-answers to the es-questions discussed above – will trade off against one another if the explanations are kept autonomous is trivial: they are explanatory virtues in view of the eg-answers being unique to one of the two theories. If one appeals to representationalism alone, one does not have access to the eg-answers to (REL-1) – (REL-3), for instance. Things get more interesting when considering the integration of the two

explanations. As things stand, integrating the two explanations (by appealing to relationalism when considering the issues involved in (REL-1) – (REL – 3), and appealing to representationalism when considering those involved in (REP-1) and (REP-2), and constructing the remainder of the account in whichever manner yields internal consistency) seems to yield a greater net total of explanatory virtues for the resulting view. Neither austere theory had access to all of these explanatory virtues for the reasons discussed in Chapter 4, yet their integration would.

I stated earlier that in determining whether relationalism and representationalism were strongly or weakly complementary explanations, the second notion of explanatory virtue – the kind possessed by an explanation in view of it explaining something *in a particular way* – would become relevant. Before illustrating how, it will be prudent to say more about these types of explanatory virtues and how they work. I gave the examples above of breadth and depth. The former is the result of a comparative approach: it unites many different phenomena, accounting for them via their commonalities; the latter is the result of a contrastive approach, and gets to the heart of how a single *explanandum* works by contrasting its mechanisms with those of other *explananda*. More simply: an explanation with greater breadth may explain more different aspects of a given phenomenon; whereas one with depth will tend to explain fewer, but will explain those that it does explain more comprehensively. I took breadth to operate by necessity over multiple eg-answered es-questions, whereas depth may operate over as few as one eg-answered es-question.

Let us return to the prospect of integrating representationalism and relationalism in the manner discussed previously: by appealing to whichever view provides satisfactory eg-answers to the es-questions that favour one theory or the other, and constructing an internally consistent view on this basis. Focusing on these es-questions, we see that integration of relationalism and representationalism makes the resulting explanation more complete, in the sense that a greater number of explanatory virtues are possessed by way of providing eg-answers to *all* of these es-questions.

As I have defined them, such integration also allows the integrated explanations to interact in a manner yielding greater breadth and depth in the resulting explanation. I have in mind here instances where the two individual accounts can genuinely interact, such as we saw when discussing (REP-2) in Chapter 4. Here, we discovered the potential for a hybrid account of why certain hallucinations and veridical perceptions were subjectively indistinguishable. The proposed view appealed simultaneously to representational content to explain the seemingly shared phenomenology, whilst utilising epistemological disjunctivism to account for the epistemological difference between the two states. Here, arguably, a greater number of experiential states are explained via the commonality of shared representational content, whilst the epistemological difference upheld by McDowellian disjunctivism allows a simultaneous deeper account to be given of their differences. In addition to these, it is trivial that by integrating the two explanations, the overall explanation offered is deeper, as by its design it explains things that each of the individual component explanations cannot. Either reading works for arguing that it is a better explanation – the net amount of explanatory virtues of both sorts increase upon integration. It is for this reason that I conclude that *relationalism and representationalism are strongly complementary explanations*.

To be clear, this conclusion involves an integration of two explanations: relationalism and representationalism. These are explanations of a single phenomenon – perception – but can still be read as constituting eg-answers to a single es-question: something like ‘what is the fundamental philosophical account of perception?’. In the following chapter, I shall make a case that the ‘fundamental’ aspect of this question makes it, in certain respects, a bad question. What is important to the present point is that this question operates at a coarser-grained level than the more particular es-questions that each theory offers eg-answers to, such as those identified in Chapter 4, which explain individual aspects of perception. It is important to interpret the conclusion given previously, therefore, as holding that relationalism and representationalism are strongly complementary

explanations *relative to this coarser-grained es-question*, and this is so in view of the heightened explanatory virtues across the board gained by integrating the explanations in a way that maps on to the eg-answers given to es-questions (REL-1) – (REL-3), (REP-1) and (REP-2).

There is one point left to address before bringing all of this together to overcome the Screening Off Argument. For all that has been said so far, relationalism and representationalism are taken to be strongly complementary in virtue of cases where one of the two theories can provide an eg-answer to an es-question that the other cannot. Yet it seems plausible that there are cases where *both* theories appear to offer distinct accounts of the same precise aspect of perception. It seems that integrating the explanations will hit a stumbling block in such cases, and so my view must provide an account of what to do in such cases. I argue that when addressing these cases, we should treat the particular explanations offered of the single aspect of perception as *weakly complementary*.

This will be possible by appealing to the distinct explanatory interests being addressed by the distinct accounts being offered. In offering distinct accounts in such cases, relationalism and representationalism offer different eg-answers to the same es-question. I argue that at a finer level of grain, they will end up being different eg-answers to more specific, *different* es-questions. If this is so, then whilst they will account for the same aspect of perception, they will do so in different ways *which reflect different explanatory interests*, as exhibited by these finer-grained es-questions. As a result of this, each will possess different explanatory virtues relative to explaining uniquely the finer-grained *explanandum*, or by explaining it in a deeper way relative to the inquirer's explanatory interest.

In this way, I conclude the following: (i) the two accounts are legitimate explanations (*within* relationalism and representationalism respectively); (ii) they are autonomous in virtue of being distinct eg-answers to what is, at the initial coarse-grained level, the same es-question; and (iii) they possess, at the finer-grained level of es-question, different explanatory virtues that will

trade off against one another relative to the explanatory interests highlighted by those distinct es-questions. Claims (i) – (iii) satisfy the conditions for explanations being weakly complementary, and so these remaining cases where relationalism and representationalism appear to offer competing explanations can be accounted for by appeal to the particular explanations in question being weakly complementary. There is no reason to privilege one such eg-answer over another, save for how well it addresses particular explanatory interests indicated by the finer-grained es-question. This too will have ramifications for the Screening Off Argument that I will discuss in the chapter that follows.

Having introduced the notions of weakly and strongly complementary explanations, and outlining what it would be for explanations to be complementary in either sense, I have argued in this chapter first that relationalism and representationalism, *qua* explanations of perception, satisfied conditions on being complementary explanations of one sort or the other, with the determining factor being whether their different explanatory virtues are magnified when the explanations are integrated. On the whole, by appeal to the es-questions (REL-1), (REL-2), (REL-3), (REP-1) and (REP-2), I have argued that this is the case and that, *qua* explanations of perception, relationalism and representationalism are strongly complementary. Finally, I addressed what to make of cases where the two views appear to offer competing eg-answers to the same es-question. Here I argued that the two theories, *qua* explanations of perception, are autonomous. However, it can also be held that these seemingly competing eg-answers, *qua* explanation of a particular aspect of perception, are actually eg-answers to distinct, finer-grained es-questions. In this respect, their explanatory virtues will trade off against one another, and so in these cases the explanations can be treated as weakly complementary, with preference being rightly allocated to the eg-answer to the es-question that more accurately reflects the explanatory interests of the inquirer. In the following chapter, I will bring the account I have proposed in this one to bear on the Screening Off Argument, and thus on the relationalism versus representationalism debate.

6: Conclusion: A pluralist reconciliation of relationalism and representationalism

The account articulated at the end of Chapter 5, drawing on all preceding chapters, holds that relationalism and representationalism are strongly complementary with respect to the coarse-grained es-question(s) that they are generally taken to address as theories of perception. I have in mind here an es-question such as ‘what is the correct philosophical account of perception?’ or something similar. I argue that instances where the two theories appear to offer different accounts were crucial to my argument that they are complementary explanations, and this is so for two reasons.

The first relates to instances highlighted by the es-questions identified in Chapter 4: (REL-1), (REL-2), (REL-3), (REP-1) and (REP-2). These instances are central to the claim that relationalism and representationalism are strongly complementary explanations because one of the conditions of being complementary in this respect is that, when integrated, the resulting explanation is a better explanation than the component explanations in some respect. I hold that the nature of these es-questions is that integrating relationalism and representationalism into a view that involves the eg-answers to each of them, as identified in Chapter 4, would satisfy this criterion. Moreover, in virtue of being the best eg-answers to the es-questions that were available, they also help to map out how relationalism and representationalism are to be integrated.

The second reason that instances where relationalism and representationalism appear to offer different accounts are important to my view involves the role that weakly complementary explanations play in the account. These instances are not like the es-questions from Chapter 4, but rather are es-questions to which both representationalism and relationalism can provide an eg-answer. The thought here is that these es-questions are finer-grained than those with respect to which the theories are strongly

complementary. Thus, they serve as explanations of more specific aspects of perception. Since they do not obviously integrate, despite being legitimate and autonomous, I argue that explanatory pluralism suggests that our means of evaluating which is to be preferred should be determined by one's explanatory interests. Drawing on material from Chapters 3 and 5, I submit that, relative to one's specific explanatory interests, the two explanations will possess different explanatory virtues that trade off against one another. Consequently, these finer-grained relationalist and representationalist explanations of more specific aspects of perception are, I conclude, weakly complementary.

As it stands, the proposed account falls short of establishing the conclusion that the Screening Off Argument can be overcome, much less the conclusion that relationalism and representationalism are compatible. Establishing these conclusions will be the central tasks of this final chapter. The structure will be as follows: I shall provide an argument that, in view of the es-questions identified in Chapter 4, the Screening Off Argument can be overcome; once the Screening Off Argument has been dealt with, I shall appeal both to the strong and weak ways in which I take relationalism and representationalism to be complementary in order to establish that relationalism and representationalism are, on this account, compatible. Thereafter, I shall address three potential objections to the view as a whole (rather than to its individual relationalist or representationalist components), highlight two features that speak in its favour, and close with some remarks on prospective avenues for future study.

6.1: *Overcoming the Screening Off Argument*

My overcoming of the Screening Off Argument involves demonstrating that Explanatory Redundancy is necessary to the success of that argument, and showing that it is not satisfied by either relationalism or

representationalism. For convenience, I state here one version of the Screening Off Argument:

The Screening Off Argument:

- 1) Suppose (for the sake of *reductio*) that perceptual experiences are representational (in the sense that representationalism says they are).
- 2) Relationalism is not explanatorily redundant with respect to what we want our philosophical theory of perception to explain (a commitment of relationalism).
- 3) The fact that perceptual experiences are representational is sufficient to explain what we want our philosophical theory of perception to explain (a commitment of austere representationalism).
- 4) If perceptual experiences are representational, relationalism is screened off as explanatorily redundant with respect to what we want our philosophical theory of perception to explain (from 3).
- 5) Relationalism is screened off as explanatorily redundant with respect to what we want our philosophical theory of perception to explain (from 1 and 4).
- 6) Contradiction (2 and 5); perceptual experiences are not representational.³⁷⁵

Recall that there are two versions of the argument, and this is just one. The idea is that the tenets of representationalism (as above) or of relationalism can be assumed to generate the *reductio* argument against the view. I argue in Chapter 5 that the inference from premise (3) to premise (4) requires appeal to the following claim:

Explanatory Redundancy: If x explains everything that needs to be explained, then y is explanatorily redundant.

Without this, there is no reason to hold the fourth premise true on the basis of the third, as the argument does. Consequently, either version of the Screening Off Argument as a whole can be rejected, provided that Explanatory Redundancy is not satisfied by either relationalism or representationalism.

³⁷⁵ Parenthetical information added to premise (3) for clarity only.

I argue that there are two readings of Explanatory Redundancy – an interest-dependent reading and an interest-independent reading. The former restricts the scope of ‘everything’ to certain explanatory interests or aspects of perception, whereas the latter has no restriction among explanatory interests salient to the phenomenon. The interest-dependent reading renders Explanatory Redundancy (at best) a false conditional, since it is by no means true that relationalism (say) explaining everything that certain explanatory interests would like to have explained entails that representationalism is explanatorily redundant. There are other perfectly legitimate aspects of perception that it may explain which relationalism cannot. Consequently, the austere theorist must establish that their theory satisfies the antecedent of the interest-independent reading of Explanatory Redundancy.

This reading *is* a true conditional, however I submit it has a false antecedent: neither relationalism nor representationalism can satisfy it. Let us address the prospects of representationalism first. That representationalism cannot explain everything we would like a philosophical theory of perception to explain, I argue, is demonstrated by three of the es-questions I identified in Chapter 4:

(REL-1): How does experience provide the perceiving subject with *conclusive* reasons for applying certain empirical concepts in judging their environment as being a particular way?,

(REL-2): What makes demonstrative reference possible?

(REL-3): How do we determine whether an erroneous experience is illusory or hallucinatory?

Of the es-questions to which relationalism can provide satisfactory eg-answers, these are selected because – as I argue in Chapter 4 – representationalism *cannot* do so. As a result (REL-1) – (REL-3) highlight aspects of perception that we would want a philosophical theory of perception to explain, yet that representationalism cannot. Representationalism therefore fails to satisfy the antecedent of Explanatory Redundancy’s interest-independent reading.

Relationalism fares no better, as Chapter 4 also identifies the following es-questions:

(REP-1): How do subjects of veridical experiences perceive the kind properties of mind-independent objects?

(REP-2): Why are certain hallucinations and veridical perceptions subjectively indistinguishable?

These constitute the counterpart to the preceding three es-questions, and although selected because representationalism can provide them with satisfactory eg-answers, they are also selected because relationalism cannot. Since they highlight aspects of perception that we would want a philosophical theory of perception to explain, and relationalism cannot provide satisfactory eg-answers, relationalism cannot satisfy the antecedent of Explanatory Redundancy's interest-independent reading either.

Taken together, these results entail that neither relationalism nor representationalism can satisfy the antecedent of Explanatory Redundancy's interest-independent reading. Hence the conditional claim comes out true, but with a false antecedent. What this means for the Screening Off Argument is that the premise (4) above cannot be successfully derived from premise (3), and so despite being valid, the argument as a whole fails due to lack of soundness. It is perhaps possible that some version of it could go through on the basis of the interest-dependent reading of explanatory redundancy, but this would be restricted in its ambition to simply the explanatory interests that one theory or the other explained everything about. Any attempt to establish that representationalism and relationalism were incompatible on such a basis would be guilty of screening off as explanatorily redundant more than it was entitled to.

6.2: *The pluralist solution*

I argue that the Screening Off Argument can be overcome by appealing to the es-questions I identified in Chapter 4. If this is right, then austere relationalists and austere representationalists cannot utilise the Screening Off Argument to support the austerity of their positions. What it does not establish is that relationalism and representationalism are compatible; it merely establishes that they cannot be argued to be incompatible in a certain way. Granted, I take the Screening Off Argument to be the most popular way of establishing their incompatibility, but what has been established so far falls short of demonstrating compatibility of the views. This is where the arguments from Chapter 5 that relationalism and representationalism are complementary explanations of perception come into play.

Following Marchionni, I discuss two ways that explanations could be complementary: strongly and weakly. Conditions common to each sort were that the explanations in question had to be legitimate explanations – i.e. they must constitute at least one eg-answer to a reasonable es-question about their target phenomenon. They must also possess different explanatory virtues. I take there to be two sorts of explanatory virtues: those an explanation possesses in view of being admirable because it explains a particular thing, and those an explanation possesses in view of being admirable because it explains something in a particular way. The former sort can be predicated as the result of a single eg-answer to an es-question, whilst the latter tend in most cases to be predicated based on several of these. I hold that relationalism and representationalism satisfy both of these conditions: each is a legitimate explanation in the sense defined, and in virtue of their ability to provide eg-answers to (REL-1) – (REL-3) and (REP-1) – (REP-2) respectively, I argue that they have different explanatory virtues too, and therefore satisfy the common conditions for being complementary explanations.

The two types of complementarity also differ. Explanations are *strongly* complementary if, when integrated into a single explanation, they form a

better explanation in the sense that the explanatory virtues involved are enhanced as a result of the integration. *Weakly* complementary explanations are autonomous: they remain unintegrated, and their explanatory virtues trade off against one another. Since representationalism and relationalism meet the conditions for complementary explanations, and I am to argue that their complementarity has bearing on whether or not they are compatible, my positive account must address what sort of complementary explanations they are.

In the latter part of Chapter 5, I argue that relationalism and representationalism are *strongly* complementary explanations. This commits them to being integrated into an explanation that is better than its individual constituent explanations in view of having enhanced explanatory virtues. I support my claim that this condition was satisfied by appealing to (REL-1) – (REL-3), (REP-1) and (REP-2). I argue that the integration of representationalism and relationalism should be drawn up in such a way that the integrated account could provide eg-answers to *all* of these es-questions.

This is possible because, as I have stated, austere relationalism is not committed to there being no perceptual representation any more than austere representationalism is committed to there being no perceptual acquaintance relation. What each is committed to is that the apparatus proposed by the other does any indispensable explanatory work in accounting for perceptual experience, and so the notion that both things obtain is not in itself precluded. By appealing to both, in the manner outlined, the proposed view integrates relationalism and representationalism in such a way as to possess more explanatory virtues, insofar as providing eg-answers to these es-questions – and therefore offering good explanations of the relevant aspects of perception – constitutes possession of explanatory virtues. The integrated explanation, by answering all of the es-questions under discussion, possesses more explanatory virtues than either of its constituent unintegrated explanations. I do not argue for the conclusion that the integrated account possesses enhanced explanatory virtues of the

second sort – such as breadth and depth – on this basis, though I do not think it implausible. Either way, the above ensures all conditions on strongly complementary explanations are met, and it is concluded that relationalism and representationalism are strongly complementary explanations.

Insofar as the proposed view involves the integration of accounts discussed, I take the question of compatibility, *relative to these es-questions*, to have been resolved. My positive account also features weakly complementary explanations, however. This is due to the relativity just emphasised. In virtue of being integrated in the way required of strongly complementary explanations, relationalism and representationalism are more specifically strongly complementary relative to the coarse-grained question that each was designed to answer. This coarse-grained question is something like ‘what is the correct philosophical account of perception?’, and the two are strongly complementary because the integrated account provides satisfactory eg-answers to es-questions that constitute an overall eg-answer to this coarse-grained es-question.

There are plausibly going to be instances, however, of es-questions to which both relationalism and representationalism can provide differing satisfactory eg-answers. Here, it looks as though the two theories will not straightforwardly integrate in the manner described above, and so compatibility of the two theories cannot be defended in the same way. In these instances, I argue, one should focus on the fact that these instances constitute specific es-questions which are of a finer grain than the coarser-grained ‘what is the correct philosophical account of perception?’ es-question. That relationalism and representationalism do not integrate neatly with respect to these es-questions is no problem for my account, however, because with respect to these finer-grained es-questions, I hold that relationalism and representationalism are *weakly* complementary.

This is possible in virtue of the explanatory pluralist methodology I have been applying throughout, according to which one explanation is not to be preferred to one another in all circumstances by necessity, but rather in view

of how well it addresses one's explanatory interests. Of course, one might discount a proposed explanation because of internal incoherence, contradiction or inconsistency, and prefer alternatives simply on the basis that they do not have these problems. The point is that when there are not inherent flaws such as these, there are no robust preference or exclusion rules on evaluating explanations. On this view an explanation is simply a statement (or set thereof), and only *becomes* an explanation in virtue of how well the statement constitutes an eg-answer to an es-question – a matter itself contextually determined by what the inquirer wishes to understand. On this account, we can identify which of two seemingly competing eg-answers is better suited to an inquirer's explanatory interests by refining the es-question until it really latches on to what they wish to understand, and then assessing which of the two (or more) eg-answers most effectively addresses this more specific es-question.

When relationalism and representationalism offer seemingly competing eg-answers to single es-questions, of the sort that make the sort of integration of the theories we saw earlier seem impossible, these instances are best addressed by utilising the method outlined in the previous paragraph. By taking these instances as cases where a specific, finer-grained es-question has been asked by some inquirer, we can take it that this inquirer possesses explanatory interests: there is something that they wish to understand which has prompted the asking of *this* es-question. By identifying what it is that they really want to understand, I argue, we can then evaluate the respective explanations offered by relationalism and representationalism of this particular aspect of perception. The one to be preferred will be the one whose explanatory virtues most closely align to the inquirer's explanatory interests. In virtue of possessing distinct explanatory virtues and not being integrated, these explanatory virtues will trade off against one another, rendering relationalism and representationalism as autonomous relative to these finer-grained es-questions, and therefore weakly complementary with respect to them.

If I am correct, then the view I am proposing incorporates both sorts of complementarity, applying them to what I take to be all of the sorts of es-questions that a philosophical theory of perception should be providing eg-answers to. That strongly complementary explanations are compatible is trivial in view of the required integration. That weakly complementary explanations are compatible I take to be true because of the purpose and nature of explanations: they exist to aid understanding and, on the view I endorse, they are statements that become explanations only because they address an inquirer's explanatory interests. Weakly complementary explanations are compatible in a context-sensitive way, in that both are appealed to, and which one is used to provide an eg-answer to a particular es-question depends upon one's explanatory interests. This context-sensitive compatibility of explanations is not problematic, I argue, precisely because explanations are themselves context sensitive.

In virtue of these arguments, I conclude that relationalism and representationalism are shown to be complementary explanations across the board (i.e. with respect to all es-questions a philosophical theory of perception might reasonably cover) and, therefore, are compatible across the board, albeit in different ways for different relevant es-questions. The relationalism versus representationalism debate itself pertained to the views as complete theories that set out to answer the question 'what is the correct philosophical account of perception?'. Regarding that debate, then, I take the two theories to be compatible in view of being strongly complementary explanations. I view them as weakly complementary only with respect to certain, more specific aspects of perception that they each appear to explain.

6.3: *Concerns*

I now address three potential concerns that I can anticipate one having about the view proposed. The first can be posed as the question: is the view being advanced a theory of perception or a meta-theory? It would seem that in

order to be a reconciliation of relationalism and representationalism, the proposed view must operate at the same level as them: i.e. it must be a theory, rather than a meta-theory. On the other hand, any independent effort to demonstrate their compatibility by refuting grounds for the opposite conclusion – such as the Screening Off Argument – would seem to require a meta-theory.

In response, I contend that the view I propose is an integration of various constituent parts of relationalism and representationalism. What my view subsumes are individual eg-answers to es-questions pertinent in the philosophy of perception, rather than the entire theories wholesale. This guarantees that the view operates at the same level as relationalism and representationalism, and therefore is reconciliatory in nature. The way in which this reconciliation was possible, however, is highly reliant on methodological considerations: the context sensitivity of explanations for its efficacy is crucial in establishing the claim that relationalism and representationalism are compatible in the manner outlined. In this respect, though not a meta-theory, the view allows a determinate stance to be taken to the question of compatibility of the two theories in view of the meta-theoretic underpinnings it has.

The second concern is summarised by the following question: to what extent is my view a reconciliation of representationalism and relationalism? By this it is meant: to what extent is the view a reconciliation of these, as opposed to simply being some alternative view? Such a concern would be motivated by the fact that the view does not appeal to either austere relationalism or austere representationalism in their entirety. Broadly, the view is not a reconciliation of either *austere* theory, insofar as it advocates the severance of the elements of austerity from both. It is this austerity – that the theory in question is the correct one *and the other is not* – that my view holds to be in fatal error. My view does neither reconciles, nor aspires to reconcile, these erroneous elements.

What it does reconcile is the respective explanatory apparatus of those theories: the notions of representation and the perceptual acquaintance relation. These, I argue, both play vital explanatory roles in the philosophy of perception. They are composed, on my view, of various eg-answers to es-questions which, in turn, are simply statements that are rightly viewed as explanations in view of a context determined by explanatory interests of the inquirer. Taking one of these statements as an example, call it S_1 , it seems fair to say that if S_1 is part of the reconciliation, there will be a statement in my account, S_2 , that corresponds to it. This will serve as an explanation in precisely the same way as it did in the austere theory it is taken from, and has the status of explanation in view of the same explanatory interests, and has the same content. For these reasons, it can be concluded that $S_1 = S_2$, and the subsuming of statements like this entail a reconciliation in the manner outlined.

Consequently, my broad answer to the concern is: it depends. More specifically, it depends on which question is being asked in the articulation of the concern. If the question is about whether the aspects of austere relationalism and austere representationalism that have explanatory utility are being reconciled: yes, the view is such a reconciliation. If the scope of the question covers reconciliation of those theories, including their elements of austerity: no, it is not.

Finally, I address a more specific concern regarding the methodological underpinning of my account. Considering the solution to the previous worry, what I am reconciling here are parts of different explanations (relationalism and representationalism) of a single phenomenon (perception). These parts are eg-answers: statements that answer es-questions. As we saw, statement S_1 from one of the austere theories has an analogue S_2 in my account. The account is a reconciliation, we saw, and so $S_1 = S_2$. Both are the same statement, with identical content, playing the same roles, and this is based on explanatory context, which is determined in turn by explanatory interests. Thus we have context sensitivity and determination explaining why $S_1 = S_2$, and explanatory interests playing the role of determining the context.

The concern runs as follows: on pain of regress, what are the identity or individuation conditions of explanatory interests? Suppose explanatory interests are context sensitive. Call whatever determines *that* context x . What are the identity or individuation conditions for x ? If something robust cannot be identified to answer this line of questions, then we have a regress.³⁷⁶

Fortunately, I think that we can halt the regress at the level of identity/individuation conditions for explanatory interests. Assuming the es-question/eg-answer framework, presumably the explanatory interests are what leads the inquirer to formulate the es-question. Presumably, then, these are to be determined by the inquirer's epistemic states: beliefs, expectations of what a good eg-answer should be, background knowledge and presuppositions.³⁷⁷ The determination also plausibly involves other intentional states of the inquirer; namely: desires, at least their desires to understand or have certain things explained. If the desire is for something robust to identify or individuate explanatory interests, it can be argued that we should look to the es-question: this is what *encodes* these notions, and does so to a reference class that suitably disambiguates what the inquirer is actually looking for from what they are not.

A coarser-grained response to the worry is that *we* (as inquirers) determine our explanatory interests, and therefore it is the inquirer's subjective individuation of their own explanatory interests that matters. This may seem arbitrary, though I do not believe it is worryingly so. We are individuating explanatory interests, not worldly things like mind-independent objects. Even perceptual experiences seem prone to individuation in a less arbitrary fashion than explanatory interests – the latter are not worldly in any robust sense and so, the thought goes, we should not expect our means of individuating them to be worldly in any robust sense either. Arguably, to the extent that they simply constitute states of mind, one can plug in one's preferred theory of individuating those – at that point it is not a problem for

³⁷⁶ I am grateful to Melissa Ebbers for raising this point.

³⁷⁷ See Schweder (1999: 116 – 117)

my view in particular, and the debate regarding which of those theories is correct is beyond the scope of the present discussion.

6.4: *Features*

Having dealt with some concerns, I now point out some positive features of my view (aside from resolving the relationalism versus representationalism debate). Firstly, if the view is correct, then proponents of several prominent theories of perception are, to some extent, mistaken. They commit the error of taking cases in which their preferred theory appears particularly strong – in virtue of it apparently explaining some aspect of perception that the allegedly competing theory seems unable to explain – as evidence that the alternative theory can be screened off as explanatorily redundant. In virtue of highlighting the methodological error being committed here, my view explains this mistake – providing a perfectly intelligible story as to how and why it occurs, whilst still making sense of the critical importance of such cases.

The austere theorists are not wrong that the cases described above – cases, I have argued, such as those highlighted by (REL-1), (REL-2), (REL-3), (REP-1) and (REP-2) – demonstrate the strength and explanatory utility of their theory. They are not even wrong that such cases demonstrate these features of their view whilst highlighting the shortcomings of the other theory. They are wrong about the significance of this. By presuming a methodology of explanatory monism, they can uphold the view that there is indeed a fundamental account to be given – a One True Theory – of perception. Then, since they successfully identify an aspect of perception that the view they take as their competition looks unable to explain, they incorporate this into the argument for their own account, generally focusing on its strengths whilst remaining silent on its weaknesses. They then conclude that the supposedly competing theory is screened off as explanatorily redundant.

My view accommodates the arguments on either side for the strengths (and, indeed, comparative weaknesses) of both theories in various areas. By applying the methodology of explanatory pluralism, it highlights not only the errors made by austere theorists on the basis of such arguments, but also provides an argument that demonstrates the erroneous nature of their austerity mentality, whilst incorporating their philosophically meritorious arguments fully.

The other feature of the view I wish to discuss relates to the fact that the integration is predicated upon reconceptualising the significance of cases where either relationalism or representationalism look particularly (uniquely) strong, and plays upon the observation that attempts to identify such cases constitutes much of the work in the debate at hand. What this allows for, I submit, is that the view I am proposing not only incorporates, but positively welcomes further examples. If it could be argued that only one of relationalism or representationalism could provide a satisfactory explanation of how we seem to perceive mind-independent objects *as* mind-independent, for example, then my view has no trouble assimilating this result into the integration of the two theories, and such cases provide further specification on the precise nature of the integration to be preferred. In this respect, the view has the potential to accommodate instances that I have not identified, yet which may be appealed to in attempts to settle the relationalism versus representationalism debate in favour of either theory.

6.5: *Future work*

The latter feature identified in the preceding section is such that it also allows for future work in the philosophy of perception to be accommodated. That is not all there is to say about directions for further study in light of the proposed view, however. I close with some brief remarks about my proposed account with a view to highlighting prospective avenues of future research.

The account I have proposed is very much focused upon, and derived from, *theories* of perception. These are taken to be explanations of a target phenomenon (namely perception) and are therefore distinct from that phenomenon. The question governing this thesis was about how these explanations interact with one another: commonly, they have been taken to be incompatible. I have argued against this claim and proposed an alternative account which integrates the two.

This is all very explanation-centric, whereas the target phenomenon of perception, I suggest, is something empirical, that takes place in the world. If this is correct, and the account offered here seems plausible, then one avenue for future research is to assess whether there is any feature of perception (i.e. as it occurs in the world) that could underwrite the view as a whole, or that could at the least utilise the insights gained from the methodology applied, even if the specifics of the resulting view are substantiated slightly differently.

Contemporary work on the unity of perception may constitute something along these lines.³⁷⁸ A view called *capacitism*, according to which relationalism and representationalism are compatible, and which holds that the austerity elements commonly applied to each are applied in error, has been proposed on the basis of perceptual capacities: capacities that function to single out mind-independent objects in the subject's environment, and which are individuated by the mind-independent objects that they function to single out.³⁷⁹ The view appears relational in nature, yet also features representational content playing an indispensable explanatory role. Consequently, the view appears reconciliatory in that "with relationalists, [it] argues that perception is constitutively relational, but with representationalists it argues that it is constitutively representational."³⁸⁰

³⁷⁸ See Schellenberg (forthcoming).

³⁷⁹ Schellenberg (2018b)

³⁸⁰ Schellenberg (2018a)

I cannot delve into the specifics of the view here for reasons of space, but it seems that perceptual capacities – which are features of the target phenomenon, rather than explanations of it – are utilised in order to draw up the view and underwrite the reconciliatory view of perception on offer. As my view proceeds in the first instance from the explanatory methodology end of the spectrum, whereas capacitism proceeds from the end of the target phenomenon, it would be interesting to see whether the two views fit well together. If they did, one could imagine an argument in favour of a more comprehensive philosophical account of perception that addresses both perceptual and methodological concerns. If they did not, one could imagine that the reasons why they failed to do so would be illuminating in themselves and would help us along the road to establishing such a view.

In addition to more philosophical features of perception, further work could also be done in addressing how well the proposed methodology squares with the findings of the empirical sciences on the topic of perception. It is worth noting that I take my view to operate at a different level of explanation to these, and so I do not take it to be in competition with them. Again, assessing how well the two sorts of take on perception may prove fruitful in our efforts to provide a unified, all-encompassing account of it.

Consequently, there are several paths for future investigation into this issue to take, at least some of which involve asking questions about perception at a broadly different level of explanation to the one I have addressed here. The area of study in question is a promising one: perception is already a topic that has seen a great deal of interdisciplinary work, so answers based on careful comparison and evaluation of the work done across multiple levels of explanation – potentially in a variety of disciplines – are not obviously out of reach, though they are beyond the scope of the present project.

Ultimately, solely in terms of being *philosophical explanations* of the phenomenon of perception, I submit that relationalism and representationalism are complementary and – however one prefers to draw that up – are therefore compatible with one another.

Bibliography

Armstrong, D. M. (1993) *A Materialist Theory of the Mind*, Revised Edition, London, Routledge.

Berkeley, G. (2012a) “An Essay towards a New Theory of Vision” in Clarke, D. M. (ed.) *Berkeley: Philosophical Writings* (2nd Ed.), Cambridge, Cambridge University Press, pp. 1 – 66.

Berkeley, G. (2012b) “A Treatise Concerning the Principles of Human Knowledge” in Clarke, D. M. (ed.) *Berkeley: Philosophical Writings* (2nd Ed.), Cambridge, Cambridge University Press, pp. 67 – 150.

Berkeley, G. (2012c) “Three Dialogues between Hylas and Philonous” in Clarke, D. M. (ed.) *Berkeley: Philosophical Writings* (2nd Ed.), Cambridge, Cambridge University Press, pp. 151 – 242.

Bird, A. & Tobin, E. (2017) “Natural Kinds” in Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy*, available online:

<https://plato.stanford.edu/entries/natural-kinds>, accessed March 2018.

Block, N. (1996) “Mental Paint and Mental Latex”, *Philosophical Issues* 7, pp. 19 – 49.

Block, N. (1990) “Inverted Earth”, *Philosophical Perspectives* 4, pp. 53 – 79.

- Bouwel, J. V. & Weber, E. (2008) "A pragmatist defense of non-relativistic explanatory pluralism in history and social science", *History and Theory* 47, pp. 168 - 182.
- Bouwel, J. V. (2004) "Explanatory pluralism in economics: against the mainstream?", *Philosophical Explorations* 7/3, pp. 299 - 315.
- Boylan, T. A. & O'Gorman, P. F. (1999) "Critical realism and economics: a causal holistic critique" in Fleetwood, S. (ed.) *Critical Realism in Economics: Development and Debate*, London, Routledge.
- Brewer, B. (2017) "The Object View of Perception", *Topoi* 36/2, pp. 215 - 227.
- Brewer, B. (2011) *Perception and its objects*, Oxford, Oxford University Press.
- Brewer, B. (2002) *Perception and Reason*, Oxford, Oxford University Press.
- Brogaard, B. (2017) "Perception Without Representation? On Travis's Argument Against the Representational View of Perception", *Topoi* 36/2, pp. 273 - 286.
- Burge, T. (1991) "Vision and Intentional Content", in Lepore, E. & van Gulick, R. (eds.) *John Searle and his Critics*, Oxford, Blackwell.
- Byrne, A. & Logue, H. (2009) *Disjunctivism: Contemporary Readings*, Cambridge, MA., MIT Press.
- Byrne, A. (2001) "Intentionalism Defended", *Philosophical Review* 110/2, pp. 199 - 240.

Campbell, J. & Cassam, Q. (2014) *Berkeley's Puzzle: What Does Experience Teach Us?*, Oxford, Oxford University Press.

Campbell, J. (2011a) "Relational vs Kantian Responses to Berkeley's Puzzle", in Roessler, J. *et al* (eds.) *Perception, Causation, and Objectivity*, Oxford, Oxford University Press.

Campbell, J. (2011b) "Consciousness and Reference" in McLaughlin, B. P. *et al*. (eds.) *The Oxford Handbook of Philosophy of Mind*, Oxford, Oxford University Press.

Campbell, J. (2005) "Transparency versus Revelation in Colour Perception" *Philosophical Topics* 33, pp. 105 - 115.

Campbell, J. (2002) *Reference and Consciousness*, Oxford, Oxford University Press.

Cavedon-Taylor, D. (2015) "Kind Properties and the Metaphysics of Perception: Towards Impure Relationalism", *Pacific Philosophical Quarterly* 96/4, pp. 487 - 509.

Chalmers, D. (2004) "The Representational Character of Experience", in Leiter, B. (ed.) *The Future for Philosophy*, Oxford, Oxford University Press, pp. 153 - 181.

Chisholm, R. (1957) *Perceiving: A Philosophical Study*, Ithaca, Cornell University Press.

Crane, T. & French, C. (2015) "The Problem of Perception", in Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy*, available online: <https://plato.stanford.edu/entries/perception-problem>, accessed January 2018.

Crane, T. (2002) "Introduction", in Crane, T. (ed.) *The Contents of Experience: Essays on Perception*, Cambridge, Cambridge University Press, pp. 1 - 17.

Dainton, B. (2011) "Time, Passage, And Immediate Experience" in Callender, C. (ed.) *The Oxford Handbook of Philosophy of Time*, Oxford, Oxford University Press.

Dainton, B. (2000) *Stream of Consciousness*, London, Routledge.

Dale, R. *et al.* (2009) "Explanatory Pluralism in Cognitive Science", *Cognitive Science* 33/2, pp. 739 - 742.

Drayson, Z. (2017) "Direct perception and the predictive mind", *Philosophical Studies*, available online: <https://doi.org/10.1007/s11098-017-0999-x>, accessed February 2018.

Evans, G. (2002) "Molyneux's Question", in Noë, A. & Thompson, E. (eds.) *Vision and Mind: Selected Readings in the Philosophy of Perception*, Cambridge, MA., MIT Press.

Fodor, J. A. (1983) *The modularity of mind: An essay on faculty psychology*, Cambridge, MA., MIT Press.

Foster, J. A. (1991) *The Immaterial Self*, London, Routledge.

- Foster, J. A. (1982) *The Case for Idealism*, London, Routledge.
- Gervais, R. (2014) "A framework for inter-level explanations: Outlines for a new explanatory pluralism", *Studies in History and Philosophy of Science Part A* 48, pp. 1 - 9.
- Gregory, R. L. (1997) *Eye and brain: The psychology of seeing* (5th Ed.), Oxford, Oxford University Press.
- Haddock, A. & Macpherson, F. (2008) *Disjunctivism: Perception, Action, Knowledge*, Oxford, Oxford University Press.
- Heck, R. (2000) "Nonconceptual Content and the Space of Reasons", *Philosophical Review* 109, pp. 483 - 523.
- Hinton, J. M. (1967) "Visual experiences", *Mind* 76, pp. 217 - 227.
- Huemer, W. (2014) "Franz Brentano", in Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy*, available online:
<https://plato.stanford.edu/entries/brentano>, accessed March 2018.
- Husserl, E. (1991) *On the Phenomenology of the Consciousness of Internal Time (1893 - 1917)*, Brough, J. B. (ed. & trans.), Dordrecht, Kluwer.
- Husserl, E. (1900/1901) *Logical Investigations 2* (2001 Reprint), Moran, D. (ed.), London, Routledge.
- Jackson, F. & Pettit, P. (1992) "In Defense of Explanatory Ecumenism", *Economics and Philosophy* 8, pp. 1 - 21.

- Jackson, F. (1982) "Epiphenomenal Qualia", *The Philosophical Quarterly* 32/127, pp. 127 - 136.
- Kellert, S. *et al.* (2006) "Introduction", in Kellert, S. *et al.* (eds.) *Scientific Pluralism*, Minneapolis, MN, University of Minnesota Press.
- Kendler, K. S. (2005) "Toward a Philosophical Structure of Psychiatry", *American Journal of Psychiatry* 162, pp. 433 - 440.
- Kitcher, P. (1981) "Explanatory unification", *Philosophy of Science* 48, pp. 251 - 281.
- Kostić, D. (2017) "Explanatory Perspectivalism: Limiting the Scope of the Hard Problem of Consciousness", *Topoi* 36/1, pp. 119 - 125.
- Kriegel, U. (2007) "Intentional Inexistence and Phenomenal Intentionality", *Philosophical Perspectives* 21, pp. 307 - 340.
- Lee, G. (2014) "Temporal Experience and the Temporal Structure of Experience" *Philosophers' Imprint* 14 pp. 1 - 21.
- Lee, G. (2006) "The Experience of Left and Right", in Hawthorne, J. & Gendler, T. S. (eds.) *Perceptual Experience*, Oxford, Oxford University Press.
- Logue, H. (2014) "Naïve Realism and Experiential Content", in Brogaard, B. (ed.) *Does Perception Have Content?*, New York, NY, Oxford University Press.
- Logue, H. (2012) "Why Naïve Realism?", *Proceedings of the Aristotelian Society* 112/2, Part 2, pp. 211 - 237.

Lyons, J. (2016) “Epistemological Problems of Perception”, in Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy*, available online: <https://plato.stanford.edu/entries/perception-episprob/>, accessed February 2018.

McCauley, R. N. & Henrich, J. (2006) “Susceptibility to the Müller-Lyer Illusion, Theory-Neutral Observation, and the Diachronic Penetrability of the Visual Input System”, *Philosophical Psychology* 19, pp. 79 – 101.

Macpherson, F. & Batty, C. (2016) “Redefining illusion and hallucination in light of new cases”, *Noûs* 26, pp. 263 – 296.

Macpherson, F. (2013) “Introduction”, in Macpherson, F. & Platchias, D. (eds.) *Hallucination*, Cambridge, MA., MIT Press.

Macpherson, F. (2012) “Cognitive Penetration of Colour Experience: Rethinking the Issue in Light of an Indirect Mechanism”, *Philosophy and Phenomenological Research* 84, pp. 24 – 62.

Macpherson, F. (2011) “Introduction”, in Hawley, K. & Macpherson, F. (eds.) *The Admissible Contents of Experience*, Chichester, Wiley-Blackwell.

Macpherson, F. (2010) “Impossible Figures”, in Goldstein, E. B. (ed.) *The SAGE encyclopedia of perception*, Thousand Oaks, CA., Sage.

Macpherson, F. (2009) “Perception, philosophical perspectives”, in Bayne, T. *et al.* (eds.) *The Oxford Companion to Consciousness*, Oxford, Oxford University Press.

Marchionni, C. (2008) “Explanatory pluralism and complementarity: From autonomy to integration”, *Philosophy of the Social Sciences* 38/3, pp. 314 – 333.

Martin, M. G. F. (2006) “On being alienated”, in Gendler, T. S. & Hawthorne, J. (eds.) *Perceptual Experience*, Oxford, Oxford University Press.

Martin, M. G. F. (2004) “The limits of self-awareness”, *Philosophical Studies* 120/1, pp. 37 – 89.

Martin, M. G. F. (2002) “The transparency of experience”, *Mind and Language* 17/4, pp. 376 – 425.

Millikan, R. G. (2009) “Biosemantics”, in McLaughlin, B. P. *et al.* (eds.) *The Oxford Handbook of the Philosophy of Mind*, Oxford, Oxford University Press.

Millikan, R. G. (1989) “Biosemantics”, *Journal of Philosophy* 86, pp. 281 – 297.

Mitchell, S. D. *et al.* (1997) “The why’s and how’s of interdisciplinarity” in Maasen, S. *et al.* (eds.) *Human by nature: Between biology and the social sciences*, Hillsdale, NJ, Erlbaum, pp. 103 – 150.

Nanay, B. (2014) “The Representationalism versus Relationalism Debate: Explanatory Contextualism about Perception”, *European Journal of Philosophy* 23/1, pp. 321 – 336.

Nudds, M. (2013) “Naïve Realism and Hallucinations”, in Macpherson, F. & Platchias, D. (eds.) *Hallucination*, Cambridge, MA, MIT Press.

- Nudds, M. (2009) "Recent Work in Perception: Naïve Realism and its Opponents", *Analysis* 69/2, pp. 334 - 346.
- Pautz, A. (2013) "The real trouble for phenomenal externalists: New empirical evidence for a brain-based theory of consciousness", in Brown, R. (ed.) *Consciousness Inside and Out: Phenomenology, Neuroscience, and the Nature of Experience*, Dordrecht, Springer.
- Peacocke, C. (1992) *A Study of Concepts*, Cambridge, MA, MIT Press.
- Phillips, I. (2010) "Perceiving temporal properties" *European Journal of Philosophy* 18, pp. 176 - 202.
- Pritchard, D. (2012) *Epistemological Disjunctivism*, Oxford, Oxford University Press.
- Prosser, S. (2016) *Experiencing Time*, Oxford, Oxford University Press.
- Prosser, S. (2011) "Affordances and Phenomenal Character in Spatial Perception", *Philosophical Review* 120/4, pp. 475 - 513.
- Putnam, H. (1975) "The Meaning of 'Meaning'", *Minnesota Studies in the Philosophy of Science* 7, pp. 215 - 271.
- Robinson, H. (2013) "The Failure of Disjunctivism to Deal with 'Hallucinations'", in Macpherson, F. & Platchias, D. (eds.) *Hallucination: Philosophy and Psychology*, Cambridge, MA., MIT Press, pp. 313 - 330.
- Robinson, H. (1994) *Perception*, Routledge, London.

Russell, B. (2001) *The Problems of Philosophy*, Oxford, Oxford University Press.

Salmon, W. (1998) *Causality and explanation*, New York, Oxford University Press.

Schellenberg, S. (forthcoming) *The Unity of Perception: Content, Consciousness, Evidence*, Oxford, Oxford University Press.

Schellenberg, S. (2018a) "Perceptual Content", *The Brains Blog*, available online: <http://philosophyofbrains.com/2018/03/06/perceptual-content.aspx>, accessed March 2018.

Schellenberg, S. (2018b) "The Foundations of Perception", *The Brains Blog*, available online: <http://philosophyofbrains.com/2018/03/06/the-foundations-of-perception.aspx>, accessed March 2018.

Schellenberg, S. (2011) "Perceptual Content Defended", *Nous* 45/4, pp. 714 - 750.

Schweder, R. (1999) "Causal Explanation and Explanatory Selection", in *Synthese* 120/1 pp. 115 - 124.

Sherman, P. W. (1988) "The Levels of Analysis", *Animal Behaviour* 36, pp. 616 - 619.

Siegel, S. (2016) "The Contents of Perception" in Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy*, available online:

<https://plato.stanford.edu/entries/perception-contents/>, accessed March 2018.

Siegel, S. (2012) “Cognitive Penetrability and Perceptual Justification”, *Noûs* 46/2, pp. 201 - 222.

Siegel, S. (2010) “Do Visual Experiences Have Contents?”, in Nanay, B. (ed.) *Perceiving the World*, Oxford, Oxford University Press.

Siegel, S. (2008) “The Epistemic Conception of Hallucination”, in Haddock, A. & Macpherson, F. (eds.) *Disjunctivism: Perception, Action, and Knowledge*, Oxford, Oxford University Press, pp. 205 - 224.

Soteriou, M. (2016) *Disjunctivism*, Oxford, Routledge.

Soteriou, M. (2014) “The Disjunctive Theory of Perception”, in Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy*, available online: <https://plato.stanford.edu/entries/perception-disjunctive>, accessed January 2018.

Soteriou, M. (2013) *The Mind's Construction*, Oxford, Oxford University Press.

Spener, M. H. (2003) *Gilding and Staining the Mind: Introspection and the Metaphysics of Visual Phenomenology*, PhD Thesis, King's College, University of London.

Thompson, B. (2010) “The Spatial Content of Experience”, *Philosophy and Phenomenological Research* 81, pp. 146 - 184.

Travis, C. S. (2004) “The Silence of the Senses”, *Mind* 113/449, pp. 57 - 94.

van Fraassen, B. C. (1988) "The Pragmatic Theory of Explanation", in Pitt, J. C. (ed.) *Theories of Explanation*, Oxford, Oxford University Press.

Vreese, L. D. *et al.* (2010) "Explanatory pluralism in the medical sciences", *Theoretical Medical Bioethics* 31/5, pp. 371 - 390.

Wright, W. (2015) "Nonconceptual Content", in Matthen, M. (ed.) *The Oxford Handbook of the Philosophy of Perception*, Oxford, Oxford University Press.