

MacGregor, Andrew Scott (2015) *A natural view of perceptual experience*. PhD thesis.

http://theses.gla.ac.uk/6361/

Copyright and moral rights for this thesis are retained by the author

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

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

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

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

Glasgow Theses Service http://theses.gla.ac.uk/ theses@gla.ac.uk

A Natural View of Perceptual Experience

Andrew Scott MacGregor

Submitted in fulfilment of the requirements for the degree of PhD in Philosophy

School of Humanities College of Arts University of Glasgow

May 2015

© Andrew MacGregor, May 2015

ABSTRACT

I offer a novel defence of radically externalist theories of perception, via a strikingly spare and broadly physicalist metaphysics. The core, motivating claim is what I call a *natural view of perception*, according to which perception involves direct awareness of our environment, such that the phenomenology of experience consists of the worldly things perceived, as they appear to the perspective of the subject. To underpin this natural view, I propose a *simple metaphysical picture* of perception, which identifies the perceptual experience with the relation of awareness holding between subject and object, a relation that can be described in familiar physical terms as a causal process involving the thing perceived and the perceiver. Distinctively, the simple metaphysical picture has no place for the notion of 'experiences' understood as distinctively 'mental' states or events internal or otherwise belonging to the subject. Although there is some limited precedent for the simple metaphysical picture of perception, I offer the first detailed argument for its role in underpinning the natural view.

The thesis offers new and detailed arguments to show that the simple metaphysical picture can not only account for normal perceptual experiences, but can also accommodate and explain other forms of sensory experience that have widely been considered to undermine the natural view of perception. These 'problem' cases include perceptual illusion, hallucination, and the role of memory and beliefs in influencing how things appear perceptually. In all of these cases, the simple metaphysical picture accounts for the phenomenology of the experience purely in terms of awareness of worldly objects, albeit in some cases objects that are not currently present in the subject's environment. The simple metaphysical picture thus promises to explain not just perceptual experience but phenomenal consciousness more generally.

The natural view is explicitly a commitment of some varieties of naïve realism, but I argue that the two theses come apart. For one thing, the simple metaphysical picture offers a solution to hallucination and other 'problem' cases quite different to the (chiefly disjunctivist) solutions offered by naïve realists. However, the most striking and novel claim advanced here is that the natural view can be defended without a commitment to realism. In this regard, I cite evidence for the subject-relativity or experience-dependence of certain perceived qualities, notably colour, and show the simple metaphysical picture allows us to square this with the natural view that colours are 'out there' in the environment.

I discuss the metaphysical implications of rejecting realism while adhering to the simple metaphysical picture, and outline a radical – and radically simple – metaphysics of the world in general that might preserve the natural view and accommodate the simple metaphysical picture of phenomenal consciousness more generally. This metaphysics takes the form of a process monism in which the governing metaphysical structuring principle is one of top-down determination, such that whole processes determine the nature of their constituent parts.

ABSTRACT			
ACKNOWLEDGEMENTS			
AUTHOR'S DECLARATION			
CHAPTER 1 - THE NATURAL VIEW OF PERCEPTION			
1.1	Introduction	. 8	
1.2	The natural view	15	
1.3	The natural view, naïve realism and relationalism	17	
1.4	A simple metaphysical picture of perceptual experience	21	
1.5	Phenomenal qualities and phenomenal character	24	
1.6	Experience as a process	28	
1.7	The limits of naïve realism	35	
CHAPTER 2 - THE CAUSAL THEORY OF PERCEPTION AND THE NATURAL VIEW			
•••••		40	
2.1	The effect view and the causal theory of perception	40	
2.2	Is the effect view a conceptual truth?	42	
2.3	Is the effect view an empirical truth?	50	
CHAPT	ER 3 - PHENOMENOLOGICAL PROBLEMS (I): ILLUSION	56	
3.1	Introduction: the arguments from illusion and hallucination	56	
3.2	Illusion and the natural view	62	
3.3	Physical illusions	64	
3.4	Cognitive illusions	67	
3.5	Optical illusions	73	
3.6	Conclusion - the argument(s) from illusion	81	
CHAPTER 4 - PHENOMENOLOGICAL PROBLEMS (II): HALLUCINATION		83	
4.1	Introduction: hallucination and the natural view	83	
4.2	The argument from hallucination	84	
4.3	Option 1a - awareness of non-worldly qualities	92	
4.4	Option 1b - uninstantiated universals	93	
4.5	Option 2 - the no phenomenology view	98	
4.6	A simple metaphysical picture of hallucination	08	
4.7	Dreams and other quasi-hallucinatory phenomena 1	15	
4.8	Missing and impossible colours 1	20	
4.9	Which constituents of the perceptual process are perceived? 1	23	
4.10	Conclusion - the argument from hallucination	27	

CHAPT AWARE	ER 5 - PHENOMENOLOGICAL PROBLEMS (III): COGNITIVE INFLUENCES OF ENESS	N .9
5.1	Introduction	9
5.2	Perceptual learning and cognitive penetration	4
5.3	Empirical evidence for cognitive contributions to phenomenology. 13	7
5.4	A simple metaphysical picture of perceptual learning	0
5.5	From perceptual learning to cognitive penetration	9
5.6	Ambiguous figures and error 15	1
5.7	Conclusion - perception as presentation and re-presentation 15	3
CHAPT OF PHI	ER 6 - PHENOMENOLOGICAL PROBLEMS (IV): THE SUBJECT-RELATIVITY	6
6.1	Introduction 15	6
6.2	Colour adaptation phenomena16	2
6.3	Intersubjective variation in unique hue perception	6
6.4	The nature of colour	′1
6.5	Subject-relativity and the simple metaphysical picture	'8
CHAPT	ER 7 - METAPHYSICAL IMPLICATIONS OF THE NATURAL VIEW 18	4
7.1	Introduction	4
7.2	Worldly objects are subject-relative	8
7.2	2.1 From experience-dependence to top-down physical determination 18	38
7.2	2.2 Wholes, processes and objects	92
7.2	2.3 The scope of top-down determination: a dilemma19	97
7.2	2.4 'Primary' versus 'secondary' qualities)0
7.2	2.5 Process monism: prospects and problems	1
7.3	The presence of the past: the temporal structure of experience 21	8
7.	3.1 What does experience tell us about time?	20
7.	3.2 The experience-dependence of the past	25
7.4	Conclusion - metaphysical implications of the natural view	8
CHAPT	ER 8 - CONCLUSIONS 23	1
REFER	ENCES	8

ACKNOWLEDGEMENTS

Thanks to staff and students of philosophy at Glasgow University and elsewhere for discussion, comments and occasional corrective looks of incredulity; and especially to Fiona Macpherson, David Bain and Alan Weir for incisive criticism and open-mindedness. Heartfelt thanks to Elaine for her forbearance and support, and to Mari and Alasdair for some questions so peculiar even I couldn't try to answer them. Thanks to my parents for keeping the kids out of my way long enough that I could have a go at my own peculiar questions. And to my grandmother, Jean Newell, for her like-mindedness.

AUTHOR'S DECLARATION

I declare that, except where explicit reference is made to the contribution of others, this dissertation is the result of my own work and has not been submitted for any other degree at the University of Glasgow or any other institution.

Andrew Scott MacGregor

CHAPTER 1 - THE NATURAL VIEW OF PERCEPTION

1.1 Introduction

When we perceive is our experience, as it seems, a straightforward confrontation with our environment, its objects and qualities? Many have thought not, including Hume, who famously derided the ordinary man for believing "the very images, presented by the senses, to be the external objects" (2007: 111). Hume thought that the notion of an "immediate intercourse between the mind and the object" was "destroyed by the slightest philosophy" (ibid.), although more recent philosophers have increasingly sought to restore credibility to some or other conception of perceptual experience as the direct engagement with the world outside our heads.

In this thesis I will explore the viability of what I will call the 'natural' view that what we call perceptual experience is just an immediate apprehension of our environment, its objects and the qualities that inhere in them.¹ A view of this sort has recent defenders amongst advocates of *naïve realism*, and I will have much to say about relevant forms of naïve realism and their success or otherwise in upholding the natural view. Insofar as we might take it already to be an implicit, default theory about the nature of perceiving, the natural view is not just a key motivation for naïve realism but also the intuitive metaphysical picture for which naïve realism seeks to provide a more robust theoretical underpinning. By contrast with theories that take a puzzling phenomenon and then propose a general theory to explain it, relevant forms of naïve realism start with an unproblematic understanding of the world and seek to conserve it in spite of seemingly problematic evidence to the contrary. That said, if naïve realism is more than an exercise in conservatism then it is because rejecting the natural view generates philosophical puzzles of its own, puzzles which naïve realism promises to resolve or avoid.

Broadly speaking, these puzzles come in two forms. First is what we might call the phenomenological puzzle of what gives rise to the qualitative or 'phenomenal' aspects of perceptual experience if not simple awareness of the world. Second are some more or less related epistemological puzzles concerning the role of perception in making it possible for

¹ By 'object' I will mean anything – any worldly entity or event – that can be thought of as bearing perceptible qualities. This might include such diverse things as items of furniture, rainbows, flashes, sounds, odours, and so on.

us to think about or refer to worldly things. It is the first sort of puzzle that I will be concerned to address in this thesis, although doing so is likely to help bring answers to the second sort into view.

As it turns out, my thesis, while defending what I call the natural view, will not prove a defence of naïve realism in any of its current forms. I will criticise existing formulations of naïve realism on two fronts. The first line of criticism has much to do with how naïve realism is expressed and in particular its adoption of, or acquiescence in, a philosophical lexicon that is ill-suited or even inimical to a proper expression of the natural view. The second line of criticism is more substantial, and concerns the inadequacy of existing naïve realist theories to explain some of the central so-called 'problems' of perception – puzzling features of sensory awareness that have led Hume and countless others to reject the natural view in the first place. These puzzling features notably include such well-known phenomena as perceptual illusions and hallucinations, as well as some less familiar phenomena like intersubjective perceptual variation.

It is in large part these phenomena that have pointed philosophers towards a view of perception as involving the occurrence of some 'inner', 'mental' states, states of the subject's mind and/or body which are caused to arise as a result of goings-on in the 'external' world. These states are what are usually called *experiences*. It is a central claim of my thesis that there are no such things as experiences in the sense just described. More precisely, my aim is to show how we can explain the nature of perception – crucially including *what it's like* to perceive – solely by appeal to the nature of the environment and our relation with it, and without recourse to peculiarly 'mental' states or properties in any substantial sense.

The widespread adherence to a conception of experiences as more or less 'inner' states of the subject is inspired in part by the thought that illusions or hallucinations can arise despite the subject's sensory engagement with her environment being tenuous or absent. The thought continues that, for all this seeming disengagement, illusions and hallucinations nonetheless seem to involve a similar sort of sensory awareness as veridical perceptual experience. Hence, in light of this subjective similarity, it is supposed that we ought to treat illusions, hallucinations and veridical perceptions as involving the subject's being in the same kind of state, i.e. enjoying an inner experience.

The conception of experience – veridical or otherwise – as 'inner' is also bolstered in the views of many philosophers and scientists by consideration of what we know about perception understood as a physical process. We know that, to perceive, a subject must be linked causally to the object perceived in a certain manner, involving the object's more or less directly stimulating some of the subject's sensory receptors and thereby causing some further neural excitation leading ultimately to a certain pattern of brain activity. Nothing in this picture obviously entails anything about the location of experience, but the 'inner' conception is again inspired by consideration of non-veridical experiences and especially hallucination. Thus, it is often noted that subjects can be induced to have hallucinatory experiences by the direct stimulation of their brain tissue, and this is taken to show that such brain activity alone suffices for the having of an experience. Since we know that brain activity is a *necessary* ingredient in the having of a genuinely perceptual experience, the seeming sufficiency of brain activity for experience in the hallucinatory case has led many to conclude that the brain activity alone is likewise sufficient for the 'experiential' aspects in the perceptual case, namely the subjective, qualitative aspects that determine 'what it's like' to have that experience. In other words, considering again the relation between perceptual *experience* and the perceptual *process*, the 'experience' part is taken to be determined or constituted by just the brain state occurring at the end of the process while the experience's being *perceptual* is constituted by the obtaining of an appropriate causal relation between that brain state and the object perceived.

This claim – that a certain kind of brain state or activity is sufficient for the occurrence of an experience, whether perceptual or not – lends credence to what is often called a *common factor view* of experience. According to this view, any set of subjectively indistinguishable experiences, whether veridical, illusory, or hallucinatory, might share a common factor, namely the experience itself (e.g. Fish 2010: 3-5; Snowdon 1981: 188 ff.; Martin 2004: 40; Child 1992: 299). This experience is, if you like, the intrinsic core of the mental state in each case – the experiences might be exactly similar intrinsically, differing only in their causal antecedents.

Further, rather intuitive support for the 'inner' conception of experience comes from consideration of more everyday phenomena such as dreams, imagination and episodic memory. As with hallucination, in each of these cases the subject seems to enjoy an experience with a more or less rich quasi-sensory phenomenology despite a lack of sensory engagement with her perceptible environment. Since these are experiences that can be enjoyed with eyes shut and ears plugged, and irrespective of our location in the world, they are conducive to the notion that they occur inside our heads in some more or less literal sense.

These, then, are the key observations that have led many philosophers to deride as 'naïve' the natural view that perception involves direct awareness of worldly objects and their qualities: first, there are observations to do with what we can know reflectively about (the diverse range of) experiences, including the possibility of experiences (illusions, hallucinations, dreams, etc.) that fail in some way to 'match' how things are in the subject's perceptible environment; and, second, empirical observations about the physical processes involved in veridical and non-veridical experiences. These observations appear to be complementary since it seems to many that we can identify a region of the perceptual process – that part taking place within the subject's brain – which is sufficient for the occurrence of these various experiences replete with sensory qualities. The crucial final step in the denunciation of the natural view is the suggestion that brain processes might be sufficient for the having of *any* such experience, including perceptual experience.

In fact there are two related ideas here. First is a general desire for unifying explanation: we can have experiences with rich sensory character (full of colours, sounds, smells, etc.) both when worldly things are perceived and also when they are not; if brain activity suffices to explain the sensory character in the latter case then we might expect it to explain it in the former case also. Second is a more specific focus on the sufficiency claim: it is sometimes suggested that the *same kind* of brain state might be involved in a perceptual experience as in a subjectively matching hallucinatory one. If that is so, and the brain state is sufficient for the sensory character of the perceptual experience then it must also be sufficient for the sensory character of the perceptual experience. As such, there would be no further part for the worldly object perceived to play in determining the character of the experience beyond its causing the right kind of brain state.

It is arguable that, until fairly recently, the natural view that was targeted by these arguments was a view that no philosopher was expected to hold. Instead, it was a view attributed to untutored common sense and the 'ordinary man', and so served as a foil for

the sophisticated theories of philosophers. The view was therefore described as 'naïve' in the ordinary, derogatory sense of the word. More recently, however, some philosophers have mounted a defence of the natural view of perception from the sophisticated objections of others, and have accordingly described themselves as naïve realists, where the naïvety is adopted as a mark of their willingness to defend common sense in spite of the sophistication of their opponents and not in ignorance of it. Inevitably, to do so they have been obliged to tackle the key questions about non-veridical experience and the physical basis of perception, and so their theories have taken on a corresponding sophistication of their own.

Naïve realism is typically defended via what is often called a *relationalist* theory of perceptual experience – one which takes perceptual experience to be a relation between the subject and the worldly object perceived. I will describe the relationalist theory in section 1.3, and in section 1.4 I will offer a straightforward way to flesh it out metaphysically – with what I will call a *simple metaphysical picture* of experience. In subsequent chapters I will show in more detail how this simple metaphysical picture can explain the nature and character of a diverse range of sensory experiences.

My purpose is, first and foremost, to defend the natural view of perception and explore the prospects of accommodating it to various aspects of our sensory experience that have been presented as problems for naïve realism. Notable among these supposed problem cases are experiences in which the way things seem to us sensorily appears not to be explicable fully, or even partly, in terms of how things are in the world before us – cases of perceptual illusion or hallucination. I will discuss these cases in detail in chapters 3 and 4 respectively, and conclude that the natural view can be defended in light of these and other supposedly problematic features of our sensory experience. More distinctively, I will further claim that these features can be explained using the same metaphysical framework applied to perception. Specifically, these problematic features can be explained by appeal to nothing more than worldly objects and the subject's sensory relation to them, a relation explicable in wholly causal, physical terms. The upshot is that what we might call the 'phenomenology' in all these experiences – the rich sensory character which determines what it's like to perceive, hallucinate, and so on – is wholly attributable to direct awareness of objects and qualities 'out there' in the world, together with the background conditions (illumination, etc.) and the subject's perspective on those objects. That this might be true

of hallucination as well as perception is likely to seem implausible at first sight, but I hope to make this more palatable in due course.

The upshot is that a metaphysical conception of awareness that is intended to deliver the natural view of perception will prove capable of accounting for the full variety of what we might call sensory experiences or phenomenal consciousness. With some of these nonperceptual forms of phenomenal consciousness – like hallucination – the proposed account will perhaps inevitably seem less than natural; after all, hallucination is an inherently 'unnatural'-seeming phenomenon. However, what is most distinctive about the current thesis is precisely that it draws on a simple, natural view of the 'good' case of perception, and uses it to explain the seemingly 'unnatural' features of the 'bad' cases like hallucination. In this respect, it inverts the usual line of argument, which cites hallucination etc. in order to show that our natural or 'naïve' view of perception cannot be correct.

While my defence of the natural view is motivated by the seeming immediacy of perceptual engagement with our environment, there is another, perhaps less obvious, feature of ordinary perceptual experience that any theory must be able to explain. This feature is experience's *generality*: not only do we seem to encounter worldly objects immediately, but we also seem to see them *as* objects of a certain kind or as possessing qualities of a certain kind. In other words perceptual experience, although seemingly immediate, is at the same time coloured by the subject's beliefs and past experiences. This is arguably a ubiquitous feature of ordinary perceptual experience, albeit one that is revealed most strikingly in certain cases of illusion involving 'aspect shifts', as in those figure/ground illusions where we 'flip' from seeing a design as a pair of faces in profile to seeing it as a vase.

Any theory of perception with a pretence to adequacy must therefore be able to account for experience's seeming immediacy without sacrificing its generality. To that end, I will argue in chapter 5 that the natural view points to a way in which we can accommodate the influence that a subject's past experiences can have on the phenomenology of current experience. Crucially, in doing so it need no more appeal to 'internally generated' phenomenal qualities than it does when explaining such phenomena as illusion and hallucination. As far as possible, I will present the natural view as a way to explain the phenomenology of a diverse range of experiences by appeal to the antecedent nature and qualities of worldly objects. It is just this explanatory approach to the character of experience that I will present as the foremost virtue of naïve realism, a virtue that in turn underpins its other claimed virtue of explaining our capacity for knowledge about particular objects and their qualities.

Despite this, I will in chapter 6 use examples from colour experience to show that the nature of some perceived qualities poses a challenge to the natural view that can only be met by denying that those qualities have their nature wholly independently of their being perceived. Such a denial is clearly tantamount to a rejection of realism about those qualities, at least on a standard notion of realism, and this is likely to impart to the argument the appearance of a *reductio* against the natural view. The obvious worry would be that positing the subject-dependence of worldly qualities points to a sort of idealism that is at odds with my stated aim of explaining perceptual experience in terms of worldly, broadly physical goings-on rather than by appeal to peculiarly 'mental' states or properties in any conventional sense. I will, however, show that we can resolve the seeming subject-dependence of the physical world, namely that worldly objects and qualities have an essentially relational nature.

This and other metaphysical implications of the natural view will be the subject of chapter 7. Meanwhile, in the remaining sections of this chapter I will lay out the core elements of the proposed theory of perceptual experience. These are two: first, a motivating claim about the seemingly world-revealing nature of perceptual experience (which I call the natural view); and, second, a metaphysical claim about perception which is deflationary about the notion of 'experiences' *qua* mental states of the subject. I will give a full account of both claims and explain the intended relationship between them. In so doing, I will also note the similarities between these claims and certain contemporary forms of naïve realism.

In subsequent chapters I will seek to defend the natural view and establish the philosophical implications of defending it. The remainder of the thesis will attempt to answer the following questions:

- <u>Chapter 2</u>: how does the natural view fit with what we know about the physical conditions of perception, *viz.* that perceiving an object requires the instantiation of certain kinds of physical-causal processes involving both the object and the physical subject?
- <u>Chapters 3 6</u>: how can we square the natural view with some features of experience accessible from a first-person perspective, and in particular seemingly deceptive experiences such as illusions and hallucinations? And how can we square the perceptual immediacy implied by the natural view with the seeming generality of experience, i.e. its capacity to present worldly things *as* things of a certain sort?
- <u>Chapter 7</u>: what does the natural view imply about the nature of the worldly objects that we perceive, and the metaphysical structure of the world as a whole?

1.2 The natural view

The natural view proves surprisingly, and ironically, rather difficult to express. To some extent we can capture it via an account of perception's *directness*. Roughly, it is the view that, when we perceive, what we are most directly aware of are just the worldly things themselves. Unfortunately, 'directness' proves a decidedly slippery concept in the philosophy of perception, taking several meanings, not all of which capture the simplicity of the natural view (for a review, see McDermid 2001). Most weakly, one might claim to uphold the directness of perception if one held that our perceiving worldly objects is not mediated by awareness of any other objects (including mental objects, like sense-data). This might be consistent with claiming that one is nonetheless most directly aware of properties of one's mental state (such as qualia, to be discussed below) rather than worldly objects. However, the sense of 'direct awareness' at play in the natural view is as straightforward and strong as one can get: it postulates no 'inner' objects of awareness

mediating perception; nor, as we shall see, does it require appeal to distinctively 'mental' states or 'experiences' in any substantial sense. Rather, perception is to be understood as a straightforward encounter between subjects and the world, with no 'added ingredients'. Quite what this entails will be explained in section 1.4 when I elaborate on a proposed *simple metaphysical picture* of perceptual experience.

For now, we might alternatively frame the natural view as a claim about *phenomenal character*, where this is customarily used to mean the qualitative aspect of awareness – the assemblage of sensible qualities (colours, sounds, smells, etc.) that constitute *what it's like* for the subject to perceive (cf. Chalmers 2002: 248). The various qualities that go together to constitute the overall phenomenal character are typically referred to as *phenomenal properties* or *phenomenal qualities*. There are some important subtleties in the employment of 'phenomenal character' and 'phenomenal qualities' to express the natural view, and I will say more about these later. Meanwhile, we can usefully, if crudely, capture the essence of the natural view by asserting that the phenomenal qualities of which we are aware are 'out there', inhering (or suffusing, emanating from, etc.) the worldly objects perceived.

<u>The natural view (first pass)</u>: in perceptual experience we are directly aware of our environment such that the phenomenal qualities that determine what it's like to perceive are 'out there' in that environment.

This definition will turn out to need some refinement, but for now it usefully sets the natural view against those theories – like sense-datum theory – that take perception to involve the awareness of objects that are 'in the mind' or mind-dependent, so that the phenomenal qualities we are directly aware of are properties of those mind-dependent objects rather than the worldly ones we take ourselves to be perceiving. Certainly, there are good grounds for resisting sense-datum theory; not least the puzzling metaphysical status of its postulated mind-dependent objects. Rejecting sense-datum theory's account of perception as indirect does not, however, commit one to the natural view. One might hold that phenomenal qualities are not properties of either worldly objects or mind-dependent objects but are instead intrinsic properties of mental states ('experiences' in the sense that I will reject). On this construal, phenomenal qualities are often called *qualia* (Shoemaker 1991; Block 2003). Since qualia, thus construed, are qualities rather than particulars, the

proponent of qualia might say that there is no *object* of which we are more directly aware than the worldly object, albeit that we are aware of it in virtue of being aware of qualia. Again, this position is clearly different from the natural view, which takes phenomenal qualities to be 'out there', instantiated in the subject's environment.

Resistance to sense-datum and qualia theories has been motivated by both phenomenological and epistemological considerations. The phenomenological considerations centre on the claim that perceptual experience is 'transparent' in the sense that it seems (typically or inevitably depending on the strength of one's transparency claim) to acquaint us with worldly objects and their qualities and not with some supposed qualities of experience itself (Harman 1990). Relevant epistemological considerations are various, but include concerns regarding the capacity of perceptual experience to give us knowledge about our environment, to ground our thoughts about particular objects, and even to enable us to form the notion of an 'external' world in the first place (see e.g. Campbell 2002b). My stated motivation for defending the natural view is phenomenological rather than epistemological, although I will remark in the conclusion on some obvious *prima facie* epistemological advantages that it offers.

1.3 The natural view, naïve realism and relationalism

It is in part a concern over the problematic phenomenological and epistemological implications of postulating sense-data or qualia that has motivated a recent surge of interest in refashioning a kind of theory – naïve realism – that was long derided as manifestly incompatible with certain observations about experience, and in particular the phenomena of illusion and hallucination.² As I will explain in chapters 3 and 4, recent advocates of naïve realism have done much work to undercut the arguments of the theory's earlier opponents and give it some plausibility.

Naïve realism is closely allied to what has been called *relationalism* about perceptual experience. Relationalism is so named on account of its central claim that perceptual experience is a relation between, and involving, the subject and the object(s) perceived.

²² For versions of naïve realism consistent with the natural claim, see e.g. Martin 2004, Fish 2009, Brewer 2011, Campbell 2002. See also Weir's 'ultra-realism' (2004).

This might seem uncontroversial at first sight, but to see why it isn't we must note a possible, and widely endorsed, distinction between *the experience* itself – an 'inner' state that in some sense holds the sensory qualities – and the link (generally taken to be causal) between that inner state and an appropriate worldly object that makes that state truly *perceptual*. On any view, *perception* is a relation between subject and object, but the relationalist's claim that perceptual *experience* is such a relation has distinctive motivations and consequences. To get a handle on these, we must first understand what is meant by the claim that perceptual experience is a relation.

What, in the first place, is the nature of the relation between subject and object? The appropriately common-sense answer is that it is a relation of *awareness* (Martin 2006: 357). Crucially, as a relation that includes its relata, the occurring of a perceptual experience entails the existence of the worldly object perceived. It is clear that, at least this far, relationalism is consistent with the natural view that what we call perceptual experience is nothing more or less than the state of affairs that is an object's being perceived by a subject.

Again, the clearest contrast is with theories that adopt a staunchly 'inner' conception of experience, such as sense-datum or qualia theories. According to such theories, the occurrence of a certain experience involves the instantiation of certain mind-dependent objects and/or qualities, but it entails nothing about the existence or otherwise of 'external' objects or qualities. Relationalism, by contrast, holds that the occurrence of an experience entails the existence of its worldly object, such that one could not have had that very experience in the absence of that object (cf. Martin 2006: 357).

The thesis is often expressed via the claim that the worldly object perceived is literally a *constituent* of the perceptual experience:

The naïve realist claims that some sensory experiences are relations to mindindependent objects. That is to say, taking experiences as episodes or events, the naïve realist supposes that some such episodes have as constituents mindindependent objects.

Martin 2006: 354

The claims that (a) experience is a relation between (and involving) the subject and worldly object, and (b) the worldly object is a constituent of the experience thus have the same force: the occurrence of the experience entails the existence of its worldly object. This notion of the worldly object's being a constituent of the experience is consistent with the natural view that the phenomenal character manifest in experience is just the sensory character of the worldly scene perceived, with its objects and their qualities. This literal reading of the constituency claim squares with what we might think of as a common-sense judgment that when we see something that something is *just there* in our field of view.

Note, however, that the constituency claim made by Martin and other naïve realists does not entail anything about the phenomenal character of our perceptual experience. Martin himself says only that "the only sense in which we can account for the role for the object of perception as a constituent of the sensory episode is acting as a necessary condition on the occurrence of the perceptual event" (2004: 57). One could, for example, allow that perceptual experience somehow takes the worldly object perceived as a constituent without one's thereby taking that constitutive role to have a bearing on phenomenal character.³

In this vein, Fish observes that relationalist naïve realism is advanced in two distinct ways, one as a claim about phenomenal character (that it is composed of worldly objects and their properties as they are arrayed before the subject in the occurrent conditions of lighting, etc.) and the other a "more ontological characterization" according to which perceived objects and their properties are constituents of the perceptual experience (2009: 16). He points out that this ontological characterisation and the one centred around phenomenal character make quite distinct claims and need not be bound together (ibid.: 17). Nonetheless, he says that the "ontological commitments…are natural concomitants of the naïve realist's claims about phenomenal character", such that the constituency claim can be seen as a way of making intelligible the claim about phenomenal character (ibid.).

³ For example, Kennedy states that "[t]he concept of phenomenal character does not feature within my statement of naïve realism's main claim, that material objects and their perceptible properties are constituent-objects of veridical experience" (2009: 590). Despite insisting that the manifest presence of particular objects contributes nothing to phenomenal character, Kennedy is anxious to retain a 'phenomenological' role for it as part of the "core conscious nature" of perceptual experience (ibid.: 600). Since phenomenal character, by his own account, "capture[s] or embod[ies] what it's like to have the experiences" (ibid.: 577), it is hard to see what 'phenomenological' or 'conscious' role can be played by particular objects besides their contribution to phenomenal character. Kennedy's motive for detaching the constituency claim from claims about phenomenal character is evidently to avoid problems allegedly arising from the possibility of hallucination. Other naïve realists have claimed to be able to account for hallucination while preserving what I have called the natural view of perception; I will consider these further in chapter 4.

Certainly, the ontological and phenomenological characterisations are often run together in the accounts of relational naïve realists: the claim that worldly objects are constituents of the experience is reflected in the claim that the phenomenal character of experience is *constituted* by the properties of the scene perceived. As Campbell states,

the qualitative character of the experience is constituted by the qualitative character of the scene perceived [...] the phenomenal character of your experience, as you look around the room, is constituted by the actual layout of the room itself: which particular objects are there, their intrinsic properties, such as colour and shape, and how they are arranged in relation to one another and to you.

Campbell 2002a: 114-116

More concisely, Brewer asserts that it is "the mind-independent direct object itself which is constitutive of this subjective character" (2006: 172). It is also said that these mind-independent objects and properties *determine* the experience's phenomenal character or, in Martin's words, "shape the contours of the subject's conscious experience" (2004: 64). In elaborating this rather metaphorical statement, Fish reasserts that the (potentially weaker) determination claim is to be read as one of constitution: "external objects and their properties…shape the contours of the subject's conscious experience by actually *being* the contours of the subject's conscious experience by actually *being* the

Such claims by naïve realists for the relational nature of experience are apt to seem very strange. The idea that worldly objects can be constituents of experiences appears, on the face of it, positively bizarre. This appearance of bizarreness reflects the tendency, noted earlier, to think of experiences as peculiarly 'mental' states that are in some sense 'inner' and private to the subject. And it is not alleviated by naïve realists' appropriation of terms like 'experience' and 'phenomenal character' and their use of them in ways that are not sufficiently sharply distinguished from their more customary philosophical use. What is needed is a clear explanation, avoiding metaphor, of what is meant by saying that objects can be 'constituents of experiences' or that the 'mind can reach out to the world' in perception. Expressing naïve realism or the natural view in this sort of way suggests some sort of overlap between two realms of being that ought to be kept apart – the realm of worldly, physical stuff and the realm of the mind, of experiences, representations, thoughts, etc. It will be evident that what is needed is nothing less than a metaphysical theory of perception that tells us what experience is (what it is 'made of') and how it relates to the world as we understand it as a realm of physical entities. Furthermore, if

such a metaphysical account is to subserve the natural view or naïve realism then it must leave the ordinary qualities of our acquaintance – the colours, sounds, smells and so on – firmly 'out there' in the world.

It is just such a metaphysical account that I will propose in the following section, one which marks no deep (or even particularly shallow) metaphysical distinction between worldly objects and our experience of them (between 'mind' and 'world' more broadly). In describing what I will call a *simple metaphysical picture* of perceptual experience, I will say some more about what we ought to mean by 'phenomenal character' and 'phenomenal qualities' when expressing the natural view. In doing so, I aim to offer a way to cash out, in rather straightforward terms, the rather obscure and metaphorical claims of naïve realists for the relational structure of experience. By no means do I claim that the simple metaphysical picture expresses what they actually *mean* by their metaphorical claims; rather, it offers one way in which those claims might be cashed out. Crucially, it is a way that offers the explanatory virtues of simplicity and parsimony, and requires appeal to nothing more or less than the world of ordinary objects that perception seems to present us with.

1.4 A simple metaphysical picture of perceptual experience

I have remarked that the natural view of perception has no place for distinctively mental states or 'experiences' in any substantial sense. This reflects in part the claimed transparency of perception, whereby experience at least *seems* to acquaint us solely with worldly things and not with some 'experience' itself (see e.g. Harman 1990).⁴ We might in any case doubt whether it is part of any pre-philosophical view that there are or are not mental states or experiences in anything like the senses employed by philosophers; these are items in our philosophical and not our natural inventory. By appealing to such things we are precisely departing from the natural view that all we need to account for our awareness of the environment is that environment itself, and our place in it.

⁴ One can of course consistently subscribe to perception's transparency while nonetheless believing that it involves the having of experiences *qua* mental states, and even that we are *in fact* acquainted with such experiences even though we don't *seem* to be.

The natural view therefore requires nothing more than that perceptual experience is the relation of awareness holding between subject and object.⁵ We might thus describe a subject S's perceptual experience of an object O as *S's being aware of O* or *S's perceiving O*. Note that we here identify the *experience* in the perceptual case with the act or event of *perceiving*. I will take it that the event of perceiving is itself describable in purely physical terms as a process in which the perceived object causes certain changes in the physiology of the subject. If we assume that this physical description omits nothing of what is involved in perception (no special 'mental' residue for example), then identifying the experience with the act or event of perceptual experience with the act or event of perceptual experience with the perceptual process.

This view is expressed in what I will call the *simple metaphysical picture* of experience:

<u>The simple metaphysical picture</u>: what we call perceptual experience is nothing more or less than the state of affairs that is the object's being perceived by the subject, or the subject's perceiving the object; a state of affairs describable in physical terms as a causal process involving, *inter alia*, the perceived object, the subject and relevant perceptual intermediaries.

Note that, while it identifies the perceptual experience with the perceptual process or state of affairs, the simple metaphysical picture says nothing about what counts as a case of perception. That is, it does not offer a *definition* of experience – it says nothing about precisely which states of affairs would fall within the ordinary concepts of seeing, hearing, etc. In other words, it neither seeks to analyse the concept of perceiving nor offers any necessary or sufficient conditions for perception having occurred. Rather, it assumes an intuitive grasp of what it is to perceive and proceeds from there to make a claim about the metaphysical structure of the episode of perceiving. Certainly, there will be a range of states of affairs whose status as episodes of perception will be questionable, for one or other reason. For example, must subjects always be conscious to perceive? Do I perceive

⁵ I will accordingly use 'perceptual experience' to denote the relation of awareness holding between subject and object, a relation which on the view proposed does not involve any distinctively 'mental' states. I will use 'experience' more generally to denote whatever is going on that suffices for our awareness (or seeming awareness) across all sorts of cases, including not only episodes of perception but also hallucination, dreaming, imagining and so on. I will indeed go on to argue that the latter, non-perceptual episodes can also be explained fully in terms of relations of awareness holding between worldly objects and the subject.

things when they are presented to me visually but I don't notice them, as when I look at the lichen-covered branch on which a camouflaged moth sits in plain sight? These and similar questions demand a conceptual analysis of perceiving that is orthogonal to my purpose in defending the natural view and the simple metaphysical picture. What matters is that, for any episode that is decreed to be one of perception, it can be explained according to the simple metaphysical picture and so without recourse to distinctively 'mental' notions.

Note also that the relationship between the natural view and the simple metaphysical picture is not one of mutual entailment. The natural view is primarily a claim about the phenomenology of awareness - a claim about where we should locate the sensory qualities - whereas the simple metaphysical picture is a claim about what perceptual experience consists of and says nothing about phenomenology as such. In this respect, the natural view and simple metaphysical picture are closely akin to the phenomenological and ontological characterisations of relationalist naïve realism described above, and similarly lack mutual entailment. One might consistently hold that phenomenal qualities are 'out there' in the environment (the natural view) without subscribing to the simple metaphysical picture (those naïve realists who persist with talk of experiences as mental states may fall into this category). Conversely, one could subscribe to the simple metaphysical picture without accepting the natural view. One might subscribe to the simple metaphysical picture while disagreeing about where in the whole perceptual process one ought to locate the phenomenal qualities. For example, Manzotti endorses the simple metaphysical picture - he calls it a 'process-oriented view of conscious perception' – but identifies an object's phenomenal qualities (its colour, say) with the whole perceptual process by which it is perceived, rather than locating the colour in the object itself (Manzotti 2008: 179). Alternatively, one might conceivably endorse the simple metaphysical picture while denying the reality of phenomenal qualities altogether.

All the same, as with the phenomenological and ontological versions of naïve realism more generally, the two theses are close bedfellows, and the simple metaphysical picture is intended to fall out of the natural view quite naturally, as it were. In explaining the phenomenology of experience, the natural view requires nothing other than the subject, her environment and the broadly physical relation between them. While the natural view therefore has no *need* for anything other than this worldly state of affairs, the simple

metaphysical picture turns this into a positive *restriction* on what perceptual experience involves.

1.5 Phenomenal qualities and phenomenal character

For all that it is not directly a claim about phenomenal character, the simple metaphysical picture helps us to get a handle on how we should use 'phenomenal character' and 'phenomenal qualities' to express the natural view. It will be worth spending a little time clarifying our use of these terms so that we can come up with a more adequate definition of the natural view.

By 'phenomenal quality' is usually meant a more or less discrete and uniform qualitative element (coloured region, sound, odour, etc.) amongst those that collectively compose the overall phenomenal character of an experience. So it might be said, for example, that a visual experience of a tomato involves one's awareness of the phenomenal qualities of roundness and redness.

In my first pass definition of the natural view, I framed it as the view that phenomenal qualities are 'out there' in the environment. We might therefore be tempted to treat 'phenomenal' as synonymous with 'sensible' such that seeing a tomato involves awareness of the redness and roundness that are the very sensible qualities of the tomato itself. However, there are cases that complicate the proposed definition. We must allow – as must any theory -a potential mismatch between how objects appear and how they are. An apple might be green, but it will look grey under red light. If 'phenomenal quality' is used to denote how the apple appears, then the apple here has the phenomenal quality of greyness. Nonetheless, there is no actual instance of greyness here – the greyness is not 'out there' in any straightforward sense. If, on the natural view, phenomenal qualities are 'out there' in the environment, we need to be able to account for this mismatch between how things are and how they appear, and we must account for it without locating the merely apparent qualities 'in the mind'. The crucial requirement, given the refusal to locate phenomenal qualities 'in the mind', is that how things appear cannot be attributed to anything other than the environment and our perspective on it. In some cases – as in the apple case above - the appearance will not be simply attributable to the sensible qualities possessed by the object, but will be a function also of prevailing background conditions;

ambient lighting, in the case described. This allows that certain phenomenal qualities are merely apparent – the seeming greyness of the apple for example – while nonetheless being wholly attributable to how things are 'out there' in the environment. As we shall see in chapter 3, this enables us to account for many cases of perceptual illusion consistently with the natural view.

'Phenomenal character', meanwhile, is used typically to denote the property of an experience that determines what it is like to have that experience (cf. Chalmers 2002: 248). Clearly, since I am using 'experience' in a non-standard sense there is likely to be some corresponding adjustment required to the meaning of 'phenomenal character'.⁶ For one thing, any talk of 'having an experience' sits awkwardly with the definition of experience presented in the simple metaphysical picture. A worldly state of affairs involving both the subject and her perceived environment is not well described as something the subject 'has'. An experience is better understood as something that involves the subject – something that she participates in. What, then, might we mean by 'phenomenal character' in this context? First of all, it is difficult to make sense of the claim that an episode of perception - an object's looking some way to a subject - might have a *character*, phenomenal or otherwise. We might talk of something – an object or a scene – as having a certain character, but less comfortably the *seeing* of something. The ordinary use of 'character' in respect of perception is to describe the overall appearance of a scene. In this use, character is a relational property of a scene -how it appears to a subject. Or perhaps we should say that it is in part a relational property of the subject's *environment*, for a scene is arguably best understood as a portion of the world characterised in relation to the perspective provided by the perceiving subject. The notion of a scene already has the relationality built in, we might say.

Character, understood as the property of a perceived environment, is thus a matter of how that environment appears or manifests itself to a subject. We can treat the adjective 'phenomenal' as supplying a contrast with other things we denote with 'character' but

⁶ The formulation 'what it's like' needs some adjustment, too. It is in any case somewhat opaque, largely because it is difficult to discern what is denoted by 'it'. Since the concern is typically said to be *what it's like for us to have an experience*, we might think that 'it' is the experience, but then why not ask simply what the experience is like? Compare the question 'What is it like to see the red tomato?' This might be paraphrased as 'What is the red tomato like visually?' This suggests that the 'it' in *what it's like* denotes the object of the experience and not the experience itself. This accords with my claim that there are no such things as 'experiences' in any substantial sense, so that talk of *what it's like* to experience really latches onto what is experienced rather than 'experience' or 'consciousness'.

which are not so essentially tied to their sensible qualities; for example, a person's character (personality), which we discover not by mere contemplation but by long acquaintance with his behaviour. While the term 'phenomenal' might not find much use in ordinary discourse, we can take it to be roughly synonymous with 'sensory' in this context, and as denoting more abstractly something we can specify with terms like 'visual', 'auditory', and so on. This enables us to align our philosophical use of 'phenomenal character' with an ordinary use, where 'visual character' might be a fancy way of denoting how something looks, and likewise, *mutatis mutandis*, for 'auditory character', and similar locutions.

Strictly, then, a proponent of the natural view should avoid talking about the phenomenal character of *experiences*. This formulation might make sense in respect of a theory according to which, for example, we are most directly aware not of worldly objects but of properties of 'experiences' (qualia in the sense of intrinsic, non-representational properties of experiential mental states - see e.g. Block 1990). Equally, it might make sense in respect of representational theories, where what it is like to perceive some worldly object is determined by the representational properties of the experience, again construed as a mental state of the subject. In such theories it is typically held that perception involves the subject's being in some mental state or enjoying a mental episode or event (the 'experience'), whose properties, intrinsic or representational, are what determine what it is like to undergo that mental state or event – determine, in other words, its phenomenal character. It should be clear that this use of 'phenomenal character' is very different from the one that I have suggested falls out of the natural view. I will use 'phenomenal character' sparingly and only to denote the sensory character of perceived objects or scenes. Where I wish to refer, neutrally, to what it's like to have an experience perceptual, hallucinatory or otherwise – I will use the term 'phenomenology', as I have above. On my account, the phenomenology of a perceptual or hallucinatory episode is constituted by the phenomenal (sensory) character of the objects or scenes perceived/hallucinated, etc.

At a second pass we may therefore offer a more robust definition of the natural view, as follows:

<u>The natural view</u>: in perceptual experience we are directly aware of our environment such that it is the environment itself that is the bearer of phenomenal character.

This revised definition allows that a scene or object possesses phenomenal character without being committed to the claim that any or all of the phenomenal qualities that appear to contribute to that character are actually instantiated in the scene. In many 'good' cases of veridical perception, how an object appears is indeed an accurate guide to how it is – the relevant phenomenal qualities are actually the sensible qualities instantiated in the object itself. When seen in broad daylight, the apple looks green because it is green. However, the natural view, as amended, allows that phenomenal qualities are sometimes merely apparent while nonetheless accounting for that appearance wholly in terms of how things are in the environment and our perspective on them.

As we have seen, the explicit use of 'phenomenal character' to denote a property of experiences (and not, as I have advocated, a relational property of worldly objects) leads naïve realists to claim that objects *constitute* this character rather than *possessing* or *bearing* it. Campbell comes closer to my refined definition of the natural view when he says in the quotation above that "the qualitative character of the experience is constituted by the qualitative character of the scene perceived" (2002a: 114-116). However, he puts out of reach the more straightforward claim of identity rather than constitution by referring first to the qualitative character *of the experience*. After all, he cannot be predicating 'qualitative character' of the same thing in both cases since, clearly, the experience and the scene are not identical. That being so, 'qualitative character' must be being used of an experience in a different way to how it is used of a scene; in the latter use, it denotes *how a scene looks (sounds, etc.)*, whereas it would be odd, to say the least, to talk of how an *experience* looks.

My purpose in redefining 'experience' and 'phenomenal character' in light of the natural view is precisely to spare that view some of the superficial implausibility that attends recent formulations of naïve realism. This implausibility results from importing key items of terminology from an area of philosophy where they have been used to express rather opposed conceptions of experience, namely conceptions of experience as somehow 'inner' and as involving objects or properties (sense data, qualia, representations, and so

on) distinct from the worldly properties perceived.⁷ The problem for naïve realism is less the terms themselves than the failure adequately to strip them of their resulting philosophical baggage. The natural view and the simple metaphysical picture are an attempt to achieve just this baggage-stripping and so make naïve realism's phenomenological and ontological claims not just intelligible but also plausible.

I should warn that the natural view, even in its revised form, will ultimately fall short of total adequacy as an account even of veridical perception, or at least it does if we take it to require that phenomenal character is wholly attributable to the subject's *current* environment, i.e. to what is presently before the subject. In chapter 5 I will consider evidence that the phenomenal character of a scene can be influenced by cognitive factors such as memory and belief, an influence that may be all but ubiquitous in normal perception. I will argue, however, that this does not undermine the claim for perception's directness, but instead highlights that even in perception we are often aware of more than what is immediately before us in our environment. Crucially, I will use the simple metaphysical picture to show that this 'extra' element of phenomenal character is nonetheless borne by worldly objects, albeit not those immediately present before the subject. Before getting to that surprising claim, we should look a bit more closely at the simple metaphysical picture.

1.6 Experience as a process

Recall that I expressed the simple metaphysical picture in two ways: firstly, by describing the episode of perceiving as a state of affairs, consistently with our ordinary talk of such episodes; and, second, by describing it in physical terms as a causal relation or process linking the object and subject in the appropriate way. This second part of this definition should not be seen as an *ad hoc* imposition of physicalism. Rather, it too should be seen as falling out of the natural view itself. That is, if we take perception at face value it presents us with nothing more or less than ourselves and the worldly scene or objects as they are arranged before us. Nothing in this state of affairs – this awareness – gives us

⁷ Witness, for example, Campbell's claim that "experience of objects…reaches all the way to the objects themselves" (2002b: 136). This usefully exemplifies the bizarre-sounding statements that naïve realists' use of philosophically-loaded terminology can lead to. Presumably, claims like these should be understood as metaphors but, absent clear admission of this or the provision of more literal alternatives, naïve realists must accept some of the blame for the incredulous stares that their claims sometimes elicit.

cause to posit any third thing or state coming between us and the objects we perceive – no mental state or experience in that sense.

If we therefore reject as superfluous any talk of experiences or mental states in this substantial sense, then we have what amounts to a form of eliminativism about mental states, and about experiences in particular. That is, our talk of 'experiences' does not pick out a class of objects, events or states, except in the sense that it denotes states of the world in which subjects are aware of objects. However, the simple metaphysical picture is certainly not eliminativist about sensible qualities – in other words, it does not deny the rich qualitative aspect of awareness, which is a crucial aspect of what we take ourselves to be talking about when we refer to experiences. What is distinctive about the natural view and the simple metaphysical picture is that they locate this rich qualitative aspect straightforwardly out in the world where it seems to be.⁸

That there is scope for a view that denies the existence of experiences *qua* distinctively mental states while admitting the reality of sensible qualities has been defended by Stoneham in support of direct or naïve realism (Stoneham 2008). According to his 'purely relational theory of perception', perceptual experience is just the obtaining of an appropriate (awareness) relation between an object and a subject (ibid.: 313). If there is a difference between this theory and my defence of the natural view it is that Stoneham describes the awareness relation as holding between "a minded thing and a physical object" (ibid.) whereas I will describe it as holding between a physical subject and a physical object, with no need to posit 'mindedness'.

A rejection of the notion of experiences as 'inner' states is also found in other naïve realists (e.g. Snowdon 1990: 124) although, as Stoneham notes, they more often suggest that perceptual experiences are 'non-inner' mental states that somehow 'reach out' to embrace their worldly objects (Stoneham 2008: 317). On the natural view, such talk of 'mental states' here is quite superfluous if our purpose is to explain the nature and

⁸ For a discussion of eliminative materialism, see Ramsey 2013. As he notes, eliminative materialism – the view that there is just physical stuff and that mental states do not exist – has sometimes been conflated with reductive physicalism – the view that mental states are just physical states of the brain (ibid.: s.1). On the view I will propose, 'mental states' – experiences in particular – are just states of the world, understood as broadly physical. This is no reduction, however: my experience when I see an object is, on this account, nothing more or less than the whole state of affairs that is my seeing the object. I describe this below as the *simple metaphysical picture* of experience.

character of awareness since we can account for this purely by reference to the nature of our environment and our (spatial, physical) perspective upon it.

A more explicitly deflationary account of the 'mental' aspect of perception is found in Weir's 'ultra-realist' theory of perception (Weir 2004). For Weir, the 'mental' state of perceptual experience consists of nothing more than the perfectly 'external', worldly state of affairs or 'situation' that is the subject's perceiving an object to be some way F (ibid: 107). The perceived objects and properties are 'bound together' in the perceptual situation by the subject's perceptual brain state, which acts as a sort of nexus for that worldly state of affairs, which otherwise would not exist (ibid.: 109). What we call 'the mind', on this view, is just a "network of situations" successively arising in this way (ibid.: 106-7).

Explicit statements of the simple metaphysical picture can be found in the writings of naïve realists such as William Child and Mark Johnston. Child suggests the following metaphysical account of perceptual experience precisely as a way that the naïve realist might reconcile the relationalist thesis with the empirical facts about perception understood as a process:

The mental state of affairs, o's looking F to S, is not a state or event at the end of a causal chain of events initiated by o; it is, rather, a (larger-sized) event or state of affairs which itself consists in the whole chain of physical events (not merely events *within S*) by which o causally affects S. The experience is the complete state of affairs, o causally affecting S. The ultimate effect in this causal state of affairs – the state or event which lies at the end of the causal chain which starts with o – is something physical in S; but that ultimate effect is neither identical with nor constitutive of the experience itself.

Child 1992: 309

Notwithstanding the reference to a 'mental state of affairs' (which we might interpret along the lines of Weir's perceptual situations), this looks much like the simple metaphysical picture as I have expressed it. Mark Johnston similarly proposes something like the simple metaphysical picture to underpin his version of naïve realism:

[T]he relation between seeing an object and the long physical process involving first the light coming from the object and then the operation of the visual system is *not* the relation between a first mental effect and a prior physical process that causes it. Seeing the object is not the next event *after the visual system operates*. Seeing the object is an event materially constituted by *the long physical process* connecting the object seen to the final state of the visual system. Seeing the

object is an event that is (as it actually turns out) constituted by a physical process that goes all the way out to the object seen.

Johnston 2004: 138-9

Johnston's description here makes no mention of 'visual experiences' as such but merely of *seeing*. As it stands, the claim that seeing – visual perception – is as a matter of fact constituted by a certain kind of process might seem fairly innocuous, and as it stands this claim is compatible with the notion that seeing nonetheless involves the occurrence of a mental state (the 'experience') at the end of this process. Johnston, however, explicitly resists the appeal to such a picture and clearly has something like the simple metaphysical picture in mind – the experience just *is* the seeing.

Johnston's account usefully applies a contrast with the more conventional picture of experience as a more or less inner mental state caused by the worldly objects perceived. This view of perceptual experience as an effect of its object I will call the *effect view*. In chapter 2 I will explore why the effect view is so widely seen as compelling and even incontrovertible and consider claims for the effect view that appeal variously to matters conceptual, empirical and phenomenological. Refuting these claims is clearly crucial to a defence of the natural view.

A distinctive feature of Child's and Johnston's statements, and which I have incorporated into my definition of the simple metaphysical picture, is the way in which experience is recast in terms of a physical process. This identification of perceptual experience with the physical process involved in perception appears in a more fully worked-out form in the work of Manzotti, who calls it a 'process-oriented' view of perception (Manzotti 2006a, 2006b, 2008, 2011). However, as noted above, Manzotti presses this metaphysical view into the service of a theory which holds phenomenal qualities to be identical to the relevant perceptual processes themselves, rather than locating them within the *parts* of those processes where they seem to reside – for example, on the surfaces of objects. From the perspective of the natural view, Manzotti's theory therefore mislocates phenomenal qualities.

In any case, if the aim of the natural view is to explain how things appear perceptually, we might wonder what purpose is served by highlighting a physical construal of the perceptual relation. After all, the perceptual process will include all sorts of elements that don't obviously contribute constitutively to the phenomenal character of what is perceived. For example, the light reflected from an apple, and the retinal excitation it subsequently causes in my eye, are not themselves seen, even if they are necessary conditions on my seeing the apple. If the purpose of the theory is just to explain why the apple looks red then an appeal to the redness of the apple, my vantage point, and the daylight conditions would seem to suffice.

One reason to offer the physical construal alongside the statement acknowledging the obtaining of a perceptual state of affairs is that it explicitly reconciles these two forms of description. This reconciliation is of a piece with the natural view's avowed purpose of explaining perceptual experience by appeal to nothing more than what we know about the physical world. As I will describe in chapter 2, there are good empirical and perhaps even conceptual grounds for thinking that an appropriate causal link between subject and object is necessary for perception to occur, so it is clear that causation plays a key role in enabling objects to play their constitutive role in experience. More generally, the physical description of the awareness relation is intended to signal that what appears to us most directly in perception are just the very same worldly objects that we measure and describe using the methods and language of science. This needn't imply that our scientific descriptions of those objects are privileged or exhaustive. There may be perceptible features of the world that are not described in a narrowly physical or scientific account of the world. For example, our use of colour language to describe objects may not align closely with a narrowly physical description of those objects (in terms of physico-chemical structure or surface spectral reflectance) but the natural view assumes that colours as ordinarily understood are nonetheless just as much out there, inhering in the objects themselves.9 I will for this reason develop the natural view within the context of what I will call a *broadly* physicalist ontology – where the use of 'broadly' is not to make room for anything not quite physical but to acknowledge that a narrowly physical description of the world may not adequately encompass all of the features of the worldly things that we can know through perception.

A second virtue of the physical construal is that it points to a route by which we can explain seemingly anomalous or 'unworldly' phenomena, such as illusion and

⁹ The failure, in some cases, of our ordinary perceptual descriptions and our narrowly physical descriptions to line up has far-reaching implications for how we must understand the world if the natural view is correct. I will explore this in some detail in chapter 7.

hallucination, experiences in which the connection of awareness to the world seems to be awry or absent. As I will explain in chapters 3 to 6, these and other seeming problem cases can be explained as more or less 'deviant' forms of awareness of worldly objects, a 'deviance' that can be made comprehensible by recasting it in terms of the physical relation between objects and subject.

It is in such 'unworldly' experiences that we might otherwise find the motivation for our customary talk of 'mental states' as somehow 'inner' phenomena. The same may be said of thoughts, memories and dreams, phenomena that we might take to be 'inner' inasmuch as they appear more or less independent of what is happening in our environment. Since they involve what is currently 'out there' in our environment only loosely if at all, and since they follow us around instead, we tend to think of them as happening 'in here' in some sense – as goings-on within our skin.¹⁰

I will argue in chapters 4 and 5, however, that memories, dreams and perhaps even thoughts – or at least their subjective, phenomenal aspects – can after all be explained without recourse to mental states or properties in any substantial sense. Just as with perceptual awareness, a full account of what it is like to imagine or remember can be given purely by reference to worldly things. If this already seems counter-intuitive and therefore 'unnatural', it is worth pointing out that what is supposed to be 'natural' is the account of ordinary perception and not necessarily the related account of our other psychological relations to worldly things. What sets hallucinations, dreams, recollections and imaginings apart from perceptions is that they do not seem to involve an encounter with things that we can literally grasp. It is the very tangibility of perceived objects that makes it natural to take perceptual experience at face value, as an immediate sensory encounter with those objects. By contrast, the intangibility of things hallucinated, dreamed, recalled, etc., seems to put them at one remove from us as subjects. Nonetheless, I will show how we can account for the 'subjective' aspects of these experiences in the same way that we can account for the qualitative character of perceptual experience, namely by identifying those experiences as (broadly physical) states of affairs involving worldly objects and subjects. My approach here will be essentially the same as that proposed by Manzotti in his

¹⁰ It is the phenomenology of these supposedly 'inner' episodes – the fact that there is often something it is like to think, remember or imagine – that generates the pressure to treat them as having special 'mental' properties or as inhabiting some special inner realm. We needn't think of thoughts, memories or imaginings as essentially exhibiting some phenomenology; even allowing that phenomenology is a mere accompaniment to (some) thoughts, memories and imaginings is enough to make the point.

'process-oriented view' (notwithstanding our differences over where to locate phenomenal qualities). In this way we can account for the seeming 'distance' between the subject and the objects hallucinated, dreamed, etc., in terms of the nature of the relevant states of affairs and without recourse to talk of distinctively 'mental' states or properties. Accordingly, I will argue that the ostensibly 'inner' character of these phenomena is quite misleading and that they can be accommodated to the natural view in essentially the same way as perceptual awareness.

Ostensibly 'unworldly' experiences like illusions, hallucinations, and so on, have therefore been seen as presenting the most glaring challenge to a natural view of perception. However, even in ordinary perceptual awareness there is good reason to believe that the phenomenology is not always attributable simply to the qualities borne by things currently before us in our environment, but can at least sometimes be influenced by our beliefs, memories and imagination. This is true notably in cases of perceptual learning and cognitive penetration of perceptual experience. Although in my characterisation of the natural view the emphasis has been on accounting for the seeming directness or immediacy of our perceptual encounters with our environment, we should therefore acknowledge the importance of accounting also for the element of generality that enters into perception – the way in which we perceive particular objects as instances of certain kinds of objects. This feature of experience we might take to underlie the evident intelligibility of the world as we perceive it – the way that it appears as an assemblage of more or less familiar objects, objects that perceptibly instantiate certain kinds and afford certain possibilities. While these features of experience – its immediacy and generality – might seem somewhat in conflict, I will show in chapter 5 that the natural view, and the simple metaphysical picture that falls from it, offer a way to account for them compatibly and, again, without appealing to any 'internal' aspects of phenomenology. As such, I will suggest that we might see the simple metaphysical picture as offering a novel way to understand the claim – typically eschewed by naïve realists – that experience represents the world as being a certain way. It might, in this way, offer a bridge between the relationalists' notion of simple awareness and the representationalists' appeal to perceptual content.

1.7 The limits of naïve realism

I have presented the natural view as an expression of the core claim of naïve realism, at least insofar as the latter theory is motivated by phenomenological concerns. And I have proposed the simple metaphysical picture as a way of cashing out naïve realists' rather metaphorical claims about the relational structure of perceptual experience and the role of objects as 'constituents' of our awareness. As noted earlier, some self-avowed naïve realists may not endorse the natural view, being concerned more with epistemological than phenomenological concerns. Nonetheless, if naïve realism does not entail the natural view, then we might at least think that the natural view entails a form of naïve realism. However, we will have cause to question this posited entailment, and indeed to hold that naïve realism is false even while the natural view might yet be true.

Before getting to the reasons for this surprising claim, we can see that even a form of naïve realism that endorses the natural view goes a step beyond it. As its name suggests, what is distinctive about naïve realism is that it combines the natural view with a further claim that I will call the *realism claim*, namely the claim that the objects and qualities of which we are aware exist and have their nature independently of their being perceived.¹¹ Let us then define the relevant form of naïve realism as the conjunction of these two claims:

<u>The natural view</u>: in perceptual experience we are directly aware of our environment such that it is the environment itself that is the bearer of phenomenal character.

and

<u>The realism claim</u>: worldly objects exist and have their perceptible qualities independently of their being perceived.

¹¹ Different forms of realism are formulated for different kinds of subject matter, e.g. mathematical objects and moral properties, but all take a similar general form, namely that the existence and nature of the relevant entities are independent of their relations to subjects. See e.g. Miller 2010 for a synopsis and discussion. The realism relevant to the current thesis concerns perceptible objects and qualities.
Taken together, the natural view and the realism claim deliver what we might see as a key explanatory virtue of naïve realism, namely its ability to explain how things appear perceptually to the subject. The explanation proposed is simple: how things appear – their phenomenal character – is determined exhaustively by the experience-independent nature of the worldly objects and qualities perceived as they are arrayed in the scene before the subject.

Let us call this the *explanatory virtue* of naïve realism:

<u>The explanatory virtue</u>: naïve realism explains what it's like to perceive solely by reference to objects and qualities in the subject's environment as they are arranged to the subject's perspective.

I take it that the realism claim is implicit in the explanatory virtue inasmuch as the direction of explanation is supposed to run from the existence and nature of the perceived objects/qualities to the phenomenology of experience. That is to say that any explanatory value is dependent on the explanans being independent of the explanandum.

The chief objections to naïve realism can be understood as attempts to show that for certain experiential phenomena – classically, perceptual illusion and hallucination - the explanatory virtue does not hold. In other words, these are supposed to be cases in which the phenomenology of experience cannot be explained solely by reference to the objects and qualities perceived by the subject. It is therefore widely claimed that the phenomenal character manifested in our awareness cannot simply be identified with the sensory character of the worldly scene itself in the way that the natural view requires. It is the alleged mismatch between how things are before the subject and how things seem perceptually that encourages some to treat phenomenal character not as a (relational) property of the perceived environment but as a (possibly intrinsic) property of an experiential mental state. The challenge to the naïve realist is thus to preserve the explanatory virtue in the face of these seemingly problematic kinds of experience and so to keep phenomenal character 'out there', as it were.

I will show in chapter 3 that naïve realists have no difficulty squaring the explanatory virtue with at least some well-known kinds of illusory experience, since these are cases in

which the phenomenal character, misleading though it may be, is nonetheless demonstrably explicable purely in terms of the scene before the subject and its constituent objects and qualities. Hallucination poses more of a challenge, since by definition it is a kind of experience in which some of the phenomenal qualities of which the subject seems to be aware need not be instantiated in the subject's perceptible environment. As a result, the possibility of hallucination has driven some of the central arguments for and against naïve realism. I will argue in chapter 4 that available naïve realist accounts of hallucination leave something to be desired, and show that the natural view points to an alternative account that, surprisingly, brings it within the scope of the explanatory virtue.

My reason for highlighting the explanatory virtue is thus to emphasise an essential feature of naïve realism, but also as a prelude to showing that it is difficult if not impossible to reconcile with certain features of perceptual experience. I will cite some key features of colour experience which cause problems for the realism claim: first, intra- and intersubjective variation in hue perception; and, second, the structure of phenomenal colour space. Each of these cases presents a challenge to any attempt to account for the phenomenal qualities of experience – their nature and variability – by reference to the qualities that worldly objects possess independently of their being perceived. It seems from such cases that we cannot identify the phenomenal qualities with subject-independent qualities borne by the worldly objects perceived, and so the explanatory virtue cannot hold.

A corollary of this is that we cannot hold true both the natural view and the realism claim, such that we must reject naïve realism as I have defined it. Given the above-noted problematic features of colour experience, such naïve realism would entail the following inconsistent triad:

- 1. Perceived (phenomenal) colours inhere, in at least some cases, in the very worldly objects perceived (the natural view).
- 2. Colours have their nature, and qualify worldly objects, independently of their being perceived (the realism claim).
- 3. Perceived (phenomenal) colour qualities are subject-relative (call this 'subject relativity').

This effectively recapitulates a familiar form of argument against naïve realism, deployed in times where it was largely invoked as a foil for what were felt to be genuinely philosophical views of perception. As I will explain in chapters 3 and 4, the argument is more commonly constructed around the phenomena of illusion and hallucination, although there are plausible accounts of these that the naïve realist can offer without obviously endangering the realism claim. The same cannot be said for the subject relativity of colour experience.

Although I will conclude from this that naïve realism is an unstable and untenable position, I will persist in defending the natural view while rejecting the realism claim. At first blush, this might seem an absurd way to resolve the inconsistent triad, and the alternative route of abandoning the natural view might seem preferable. Certainly, it has been preferred by most philosophers when confronted by a similarly-structured problem based on the possibility of illusion or hallucination. Nonetheless, I think that the natural view remains worth defending precisely because it is such a statement of the obvious, by which I mean it is so well grounded in a compelling pre-philosophical grasp of what perception involves. One might object that the realism claim is likewise a statement of what seems obvious, so that there is an inconsistency in denying a corresponding defence of that claim. However, while the realism claim might seem to capture a common-sense view (albeit one not always endorsed by philosophers, notably including some idealists), it is arguably a more philosophical claim than the pre-philosophical insight I have tried to express in the natural view. No doubt there are common-sense insights in the vicinity of the realism claim, although these arguably have more to do with notions of causal independence, namely that we cannot change things merely by perceiving them. The notion of constitutive independence is less easily discerned in what we might imagine is a common-sense, pre-philosophical view of perception.

In any case, while the realism claim and subject relativity are clearly incompatible (the latter is the denial of the former, at least in respect of certain perceptible qualities), no such logical incompatibility comes between subject relativity and the natural view. Of course, to hold both of these claims would be to commit oneself to the claim that the inherent qualities of worldly objects (colours, etc.) are subject-relative. This, for one thing, requires a rather loose reading of 'inherent'; specifically one that does not treat it as a synonym for

'intrinsic'. More worryingly perhaps, it looks very like a kind of idealism, namely the claim that worldly objects are mind-dependent. However, I will argue in chapter 7 that the natural view opens the way to a metaphysical picture that can account for the falsity of the realism claim without succumbing to what have been seen as the more unacceptable implications of idealism. Most obviously, my denial of 'mental states' or 'experiences' construed as distinct entities, states or events occurring within the subject means that the relevant notion of subject-dependence is not to be read as *mind*-dependence or understood in terms of the ontological priority of the mental over the physical. Instead, I will argue in chapter 7 that it can be cashed out in terms of a certain kind of top-down metaphysical priority relation holding between wholes and their parts. The crucial example of this metaphysical priority for my purposes is the priority of the whole perceptual process (involving and linking object and subject) over the particular worldly objects and qualities that form parts of that process.

Notwithstanding the metaphysical contortions necessitated by subject relativity, I will show that the simple metaphysical picture allows us to explain many hallucinatory and other non-perceptual experiences in a way that *preserves* realism and hence naïve realism's explanatory virtue. As such, I will argue that the natural view and simple metaphysical picture offer an advance on some existing naïve realist attempts to account for the phenomenology of non-perceptual experience. Ultimately, however, we will be left with some aspects of sensory experience that are also explicable in terms of the simple metaphysical picture, but only at the cost of abandoning realism. Quite what this abandonment amounts to will be discussed in chapter 7.

CHAPTER 2 - THE CAUSAL THEORY OF PERCEPTION AND THE NATURAL VIEW

I have argued that the natural view is suggestive of a simple metaphysical picture of perceptual experience, namely that what we call 'experience' is nothing other than the state of affairs that is the object's being perceived by the subject. In this chapter I will make an initial survey of the prospects for the simple metaphysical picture. To that end, I will critically appraise some *prima facie* reasons for thinking that perception must involve some sort of more or less inner 'experience' – a distinctively 'mental' state or event that is caused to arise in the subject by the object perceived.

2.1 The effect view and the causal theory of perception

I have suggested that the natural view and the simple metaphysical picture might make intelligible the relationalist claim offered by naïve realists to explain the directness of perception, namely the claim that the worldly objects perceived are literally constituents of the perceptual experience. The simple metaphysical picture makes good this claim by identifying what we call 'experience' with just the worldly process or state of affairs that is the object's being perceived by the subject. Nothing metaphysically extravagant is intended by this claim. Quite the reverse: it is intended to be deflationary about 'experiences' insofar as it denies that talk of 'having a perceptual experience' is to be understood as entailing the occurrence of something – an experience – that is in some way distinctively mental and separate from the thing perceived. 'Having a perceptual experience' is, given the simple metaphysical picture, taken instead to be a convoluted synonym of 'perceiving'.¹² Otherwise, the claim that perceptual experience takes the worldly object as a constituent might strike one as implausible or bizarre insofar as it contradicts a commitment that is, explicitly or not, common to the vast majority of theories of perceptual experience. The commitment in question is a broad claim about the relationship between perceptual experience and the physical process involved in perception, namely that perceptual experience is an event or state within the subject which occurs at the end of a causal chain originating externally to the mind, such that the

¹² One might think that it is possible to perceive unconsciously, while *having a perceptual experience* expressly implicates *conscious* perceiving. In that case, we might prefer to treat 'having a perceptual experience' as synonymous with 'consciously perceiving', but without taking this to require that the conscious aspect of it involves the occurrence of an 'experience' in the substantial sense that I reject.

experience itself is an effect of the worldly thing perceived. I will call this the *effect view* of perceptual experience:

<u>The effect view</u>: perceptual experience is a state of, or event within, the subject which is caused by the worldly object perceived.¹³

The effect view is widely seen as uncontroversial and indeed obviously true. As a result, it is rarely defended as a philosophical claim in its own right.¹⁴ A philosophical claim it is, nonetheless. Doubtless what has obscured its contestable status as a philosophical claim is a failure to conceive of any plausible alternatives. And the likely worry is that a rejection of the effect view would entail detaching experience from the realm of causes altogether, an outcome with unfortunate consequences for all of the philosophical projects, metaphysical and epistemological alike, that centre on reconciling mind and the physical world.

It should be clear, however, that it is just this reconciliation between mind and world that the natural view is intended to bring about. More precisely, it follows from the natural view that the felt *need* for reconciliation is really an artefact of the effect view's conception of 'mind' as something more or less distinct from the world. As such, we can go further than the weak requirement that the physical facts about the perceptual process do not rule out the natural view or otherwise favour the effect view, and assert more strongly that the natural view is more obviously consistent with a broadly physical ontology. The purpose of the current chapter is to show that the weaker requirement is met – that the natural view

¹³ Most theories of perceptual experience will draw a distinction between the experience itself and its 'content' – what the experience is about, the state of affairs that would make the experience accurate or true. This distinction is often cashed out as holding between an inner state of the subject and what that state represents; between a representational 'vehicle' and the representational content (e.g. Dretske 2003: 68). Clearly, on this view, where the content of the experience refers to an object in the subject's environment the experience's content is not a state of or event within the subject caused by that object. In respect of such theories the effect view therefore concerns the experience itself – the vehicle rather than the content. On the view that I will defend there is no such distinction: the objects that a perceptual experience is about are constituents of the experience itself.

¹⁴ For example, Valberg states that what he calls 'the causal picture of experience' – crucially incorporating the claim that experience is the 'upshot', 'result' or 'culmination' of the familiar causal process of vision involving light, eyes, nervous excitation, etc. – is 'not in any sense a 'philosophical' view or theory'' (1992: 24). A little more circumspectly, Strawson remarks that "with the distinction between independently existing objects and perceptual awareness of objects we already have the general notion of causal dependence of the latter on the former, even if this is not a matter to which we give much reflective attention in our pre-theoretical days'' (2011: 136). We might think the latter caveat somewhat undermines his earlier claim that the stated distinction "is as firmly a part of our pre-theoretical scheme as is our taking ourselves, in general, to be immediately aware of those objects'' (ibid.: 135).

is consistent with the physical facts about perception. The wider metaphysical implications of the simple metaphysical picture will be the subject of chapter 7.

I will argue that we can reject the effect view by denying the need to postulate a perceptual *experience* that is taken to be a more or less discrete, datable event or entity occupying a particular place within the physical-causal perceptual process. As we will see later, however, this denial requires that the burden of explaining any peculiar features of perception, and of sensory experience more generally, falls back on to the world. In other words, any seemingly anomalous features of how things appear to us sensorily must, after all, prove explicable in terms of the nature of the physical world. Accordingly, in chapters 3 to 6, I will argue that the natural view is consistent with what we know about experience from the first-person perspective, and in particular that it can accommodate phenomena like perceptual illusion and hallucination which have widely been seen as problematic for naïve realism.

In the meantime, it might be thought that the effect view is supported by empirical findings regarding the perceptual process. One might furthermore hold it to be a requirement of our ordinary concept of perception. In other words, the effect view might be considered as either or both an empirical and a conceptual truth. In this section I will argue that it is neither. Clearly, the claim that the effect view is a conceptual truth is a stronger claim than the empirical claim, as the former would entail the latter but not *vice versa*. For that reason it is the conceptual claim that I will assess first.

2.2 Is the effect view a conceptual truth?

One reason that might be offered for endorsing the effect view is that the very concept of perception is a causal concept, and specifically that it is part of our ordinary concept of perception that perceiving involves the having of an experience that is caused by the worldly object perceived. This widely-held view has been called *the causal theory of perception* (hereafter, 'the causal theory') (see e.g. Grice 1961, Pears 1976).

The natural view plainly conflicts with the causal theory. This is a consequence of the natural view's deflationary interpretation of talk of 'having experiences', namely that such talk is a misleadingly complicated and metaphorical way of referring to *experiencing*

(perceiving, etc.). Certainly, the natural view allows that perception is a causal process – something made explicit in the simple metaphysical picture – but not that the causal relata are the object perceived and an experience. Rather, the relevant relata are the object perceived and some physical state of the subject. The *experience*, according to the simple metaphysical picture, is just the perceiving, otherwise describable as the whole perceptual process involving object and subject.

In this section I will argue that claims for the conceptual truth of the causal theory (and thereby the effect view) are unfounded. It will not be necessary to my defence of the natural view to show that the effect view is conceptually false, or that the natural view is conceptually true; rather, it will suffice to show that the effect view is not conceptually true. Equally, note that the causal theory is a specific form of argument for the effect view – one which seeks to establish the truth of the effect view on conceptual grounds alone. The failure of the causal theory does not, in itself, entail the falsity of the effect view. I will address alternative, empirical arguments for the effect view later in this chapter.

The causal theory can be stated as follows, where '*S*' denotes a subject and '*o*' an object (cf. Snowdon 1990: 123):

<u>The causal theory</u>: It is a conceptual truth that if *S* perceives *o*, *S* is undergoing a perceptual experience caused by *o*.

Crucially, 'perceptual experience' is not being used here to denote the state of affairs that is o's being perceived by S, as the natural view would claim. Two considerations make this clear: firstly, the experience is said to be *caused* by o, which plausibly makes sense only if o is distinct from the experience; second, even if we could make sense of o's causing a state of affairs of which it is a constituent, that would reduce the claim to tautology rather than a conceptual truth. The notion of perceptual experience being employed by causal theorists is evidently one that treats it as a more or less 'inner' state of or event within the subject.

What distinguishes the causal theory from the effect view is its presentation as a conceptual truth - a requirement of our ordinary concept of perception - rather than as a mere empirical claim or statement of fact. The causal theory is, to that extent, intended as

a form of conceptual analysis. Quite what is meant by 'conceptual analysis' in this kind of context requires some elucidation, however. In the broadest terms, we may take it to be a process by which we uncover the ways in which a concept is used, the scope and limits of its application. Insofar as it is intended to reveal something of our ordinary concepts, it is reasonable to suppose that this analysis should not rely on matters of fact that are likely to outstrip the knowledge possessed by ordinary users of the concept. *Seeing*, for example, is a concept applied in everyday situations by subjects with little or no knowledge of the physical processes involved in our visually apprehending worldly objects, so such knowledge cannot play a role in explicating the conditions and constraints on our use of the concept.

As such, when Snowdon says that "it is wrong to incorporate as an element in the analysis of a concept C any condition F which can be revealed as a necessary and essential condition for the correct application of C only by arguments relying on what are, broadly, empirical considerations", he means that we must only incorporate conditions that can be inferred from the judgments of ordinary subjects in respect of the scope of applicability of the concept concerned (Snowdon 1990: 121-2). That said, perception is of its very nature a process of interaction with the world and the source of much of our knowledge, and our mastery of the concept of perception or seeing is likewise the outcome of such interactions, so what is intended is far from an *a priori* analysis of the meaning of terms like 'perceiving' or 'seeing'.

Snowdon finds in the causal theory (as applied to visual perception) three distinct claims (Snowdon 1981: 175-6):

- necessarily, if a subject S sees an object O, then O causally affects S ("the causal thesis");
- the effect of O on S is a state of S which can be reported by a *looks*-statement,
 i.e. one of the form 'It looks to S as if...' ("the effect thesis"); and
- the causal thesis and the effect thesis are essential to our ordinary concept of visual perception ("the conceptual thesis").

In what follows, it will be important to bear in mind the difference between the causal *theory* and the causal *thesis*. The causal thesis (being one component of the overall causal theory) describes a necessary condition of *S*'s perceiving O – namely that O affects *S*. It should be obvious that it does not describe a *sufficient* condition of *S*'s perceiving O, since the object may affect the subject by some route other than via her sensory apparatus, and in a manner that does not render the object perceptible. To address this insufficiency, the effect thesis claims that the effect of O on S is to bring about a state that can be characterized in terms of how things look to S; in other words a perceptual experience. The conceptual thesis simply adds that the causal and effect theses are true by virtue of the very concept of seeing.

For my purposes, this analysis of the causal theory usefully teases apart elements that are inimical to the simple metaphysical picture from those that are not. Specifically, the simple metaphysical picture is at odds with the effect thesis, inasmuch as the latter takes the state of *S* in question to be the perceptual experience. The causal thesis, meanwhile, is something that both the effect view and the simple metaphysical picture explicitly endorse. It will therefore be quite consistent with the simple metaphysical picture should the causal thesis prove to be a conceptual truth. Of course the conceptual truth of the causal thesis would not favour one of the effect view and the simple metaphysical picture over the other, for what is at issue is the relationship between the relevant effect on the subject and the perceptual experience. By the same token, neither the effect view nor the simple metaphysical picture would be disproved if the causal thesis turned out not to be conceptually true, for there might be empirical evidence in its favour instead. (Both theories would be in trouble if the causal thesis proved either conceptually or empirically false, but this does not appear likely.)

Crucially, the principal arguments raised in favour of the causal theory are ones that support only the causal thesis. The causal thesis is typically brought out by consideration of the kinds of 'problem' cases noted above, involving various hypothetical disruptions to the causal process that, in straightforward cases of veridical perception, hold between object and subject. Taking vision as the paradigmatic mode of perception, a noncontentious case of veridical visual perception (a case of seeing) would be one where the following three circumstances obtain: (i) it looks to the subject as if there is an object of type F at location L; (ii) there is an object of type F at location L; and (iii) the object at L causally affects the subject in the appropriate manner via the subject's visual system. The kinds of problem cases that are typically used to defend the causal theory are what we might call *veridical illusions* or *veridical hallucinations* – cases in which circumstances (i) and (ii) obtain but (iii) does not – there is an F at L, it looks as though there is an F at L, but the F at L is not causally affecting the subject (Snowdon 1981: 180-1).

Broadly two kinds of case might be contrived, differing from each other in the deviance (or seeming non-existence) of the causal chain leading to an effect on the subject (cf. Snowdon 1981: 181):

- 1. *Veridical illusion*. Here a subject is standing in front of, say, a stuffed parrot and it is to her as she were seeing a stuffed parrot. Unbeknownst to her, however, there is a mirror between her and the stuffed parrot and the mirror is reflecting a matching cardboard cut-out of a parrot situated off to the side.
- 2. *Veridical hallucination*. In this kind of case the subject might again be standing in front of the stuffed parrot and it is to her as she were seeing a stuffed parrot. But now the parrot is unobstructed in her line of sight, yet she is having a hallucination that prevents her sensory apparatus registering the actual parrot but which just happens to match the character of the experience she would otherwise have of the actual parrot.

The conceptual truth of the causal thesis depends on our intuition that neither of these cases is one in which the subject sees the target object (the real stuffed parrot) despite the fact that conditions (i) and (ii) above hold, i.e. the target object is present and the subject appears to see something appropriate to the target object's being there. What is missing from each case is of course the appropriate causal connection between the stuffed parrot and the subject.

The hypothetical problem cases are the main source of justification for the causal theory. However, it should be obvious that they support only the causal thesis and not the effect thesis. What they show is that we ordinarily consider a case to be one of perception only if a causal connection holds between the subject and the object purportedly perceived. But of course the causal thesis omits any claim regarding the nature of the effect induced in

the subject. This specification is the role of the effect thesis, which effectively claims that *O* causes a perceptual experience in *S*.

That the causal thesis is *in fact* true is not challenged by Snowdon, although he questions whether its truth follows also from the very concept of vision (1981: 176 ff.).¹⁵ As noted above, whether the causal thesis is conceptually true in addition to being empirically true is of no consequence for the debate between the effect view and the simple metaphysical picture. This debate hangs entirely on the truth or otherwise of the effect thesis.

Unfortunately, in his (1981) Snowdon offers little explicit criticism of the effect thesis nor any explanation of why someone might hold it to be true. Indeed, having defined it, Snowdon gives it scarcely another mention.¹⁶ His position on the effect thesis is clearer from his (1990) despite the fact that here he addresses the causal theory in its entirety without dividing it into its constituent theses, as follows: "it is a conceptual requirement that, necessarily, if P (a subject) sees O (an object) then O is causally responsible for an experience (call it E) undergone by (or had by) P" (1990: 123). It is clear, however, that the chief target is now the alleged conceptual truth of the effect thesis: the claim that it is a conceptual truth that P's seeing O involves P's having an experience (1990: 124). Specifically, he seeks to undermine the claim that, conceptually, seeing something involves a subject's having an 'inner' experience, i.e. "an occurrence in P which is an experience and which is quite separate from O" (1990: 127).

⁵ As I will note below, Snowdon argues that if perception were a causal concept there would have to be a manifest effect end in perceptual experience, and this manifest effect end is at least not obvious given the familiar transparency of normal experience (1990: 136-7). Hyman offers an alternative objection to the causal thesis, suggesting that the illustrative 'problem' cases (i.e. veridical illusion and veridical hallucination) show not that causation is involved in seeing, but merely that it is possible causally to *prevent* one from seeing (Hyman 1992: 280). In this regard he is at odds with Child, who claims that the causal thesis is supported by our intuitive grasp of the possible ways in which our vision of an object can be *blocked* (as by having our eyes closed, or by some other object obstructing our line of sight), blocking being a causal notion (Child 1992: 311). Hyman suggests instead that X's blocking our view of Y involves the denial of an opportunity to see Y rather than *X's causing us not to see Y* (Hyman 1994: 374-5). To which Child retorts that opportunity conditions (and their denial) can be either causal or non-causal (where the latter involves merely logical entailment, such as one's being an only child prevents one from being an uncle, to use Child's own example) and that the kind invoked by Hyman is plainly the former (Child 1994: 363-4).

¹⁶ It reappears briefly as the 'looks-thesis', "namely, the claim that if S sees O then S is in an L-state", where he clarifies that his admission of such L-states is not intended to imply any ontological commitment but merely stands in for the claim that "if S sees O then it is true to say of S that it looks to him to be some way" (1981: 180 & fn. 8). However, such admission of L-states remains at least misleading insofar as they are allowed to be caused by perceived worldly objects.

Snowdon acknowledges why someone might hold the effect thesis, namely because of the possibility of 'total hallucination' - experiences that are wholly hallucinated. Here, he accepts that such hallucinations must be 'inner' experiences for which the intrinsic state of the subject suffices, since they can be had in the absence of an appropriate object (1990: 128). The effect thesis then follows via the assumption that the same kind of experience is had in the perceptual case as in the hallucinatory case, i.e. both are inner experiences in the sense defined (1990: 129). (We have here the outlines of the 'argument from hallucination' against naïve realism, an argument that I will discuss later in chapter 4, where I will there dispute the assumption that hallucinations are inner experiences in Snowdon's sense. However, this has no bearing on my purpose in this section, namely to show that the effect thesis – and so the effect view – is not a conceptual truth.) The picture being assumed is made clearer by considering a pair of possible experiences – one perceptual, the other hallucinatory – which appear the same to the subject. It is claimed that both experiences are intrinsically the same - it is the same kind of experience happening in both cases – and that what makes one perceptual is merely the appropriate causal relation to the object perceived.

As Snowdon argues, however, the assumption that perception and hallucination must involve the same kind of experience is in need of considerable support. Even if it has a marked explanatory virtue, it does not have the status of a conceptual truth. Snowdon demonstrates this by considering a 'radical alternative' to the picture assumed by the causal theory. According to this alternative picture, a statement about how things look to the subject could be made true by different kinds of experiences in the perceptual and hallucinatory cases. In the perceptual case, the experience has the object as an 'ingredient' (i.e. the object is a constituent as per the relationalist picture described in chapter 1), while in the hallucinatory case the experience has some other metaphysical structure in virtue of which the looks-statement is true (Snowdon 1990: 129). This claim - that the same looksstatement could be made true by different kinds of experience in the perceptual and hallucinatory cases – amounts to what is known as *disjunctivism* about perceptual experience; 'disjunctivism' because it posits a potentially disjunctive analysis of the looksstatement. In Snowdon's formulation this disjunctive analysis has it that a statement of the sort "It looks to S as if there is something which is F" might be made true of either of two states of affairs: (a) there actually *being* something which looks to S to be F, or (b) its merely *looking* to S as if there were something that is F (Snowdon 1981: 185). In other

words, an experience in which it seems as though things are some way might actually involve a genuine perceptual encounter with an object (might be 'object-involving') or it might not.

The crux of Snowdon's argument is that the disjunctive analysis, while it might not in fact be true, is at least not conceptually false – there is nothing in our ordinary concept of perception that rules it out (Snowdon 1990: 130-1; see also Snowdon 1981: 191). What underpins the effect thesis is an explanatory hypothesis concerning the possibility of subjectively indistinguishable perceptual and hallucinatory experiences; namely, that this indistinguishability is explained by these experiences' being metaphysically of the same kind. This urge towards a single, unified explanation of experiences in general is understandable but it cannot offer conceptual grounds for endorsing the effect thesis (and so the effect view). The most it can offer is an inference to the best explanation.

It is interesting to note that Snowdon's argument against the conceptual truth of the *causal* thesis appeals to a feature of experience that is even more obviously damaging for the effect thesis. He argues that when we perceive something "there is nothing in the occurrence which is both manifest to us and can count as an effect induced by, and hence separate from, the item seen" (1990: 136). That being so, he finds it implausible that the concept of seeing "is a causal concept with a separable experience required as the effect end" (ibid.). Snowdon is here appealing to the transparency of experience, i.e. the claim that in visual experience all we are aware of are properties of the object seen and not properties of the experience itself (ibid.: 136-7). As he puts it,

in perception there is nothing to latch on to other than the world; in particular, there is no such thing as a state produced in us, and which is manifestly distinct from the world, to which we can attend. This...is precisely how experience strikes us.

Snowdon 1990: 137

The suggestion here is that transparency deprives us of introspective evidence for the occurrence of experiences, insofar as these are taken to be distinct from the worldly objects experienced. As it stands, this is most directly a challenge to the conceptual truth of the effect thesis, and as such it provides a positive argument to bolster his conclusion that an alternative disjunctive account is not conceptually false. Thus, if the effect thesis were conceptually true, we might expect that we should be aware of the experience itself as

distinct from what it is an experience *of*. This is at least not obviously the case, and commitment to a strong transparency claim rules it out altogether. Even if we allow a weaker transparency claim, such as that we are *normally* only aware of properties of experienced objects (although we can with some effort introspect the experience itself), this would leave awareness of the experience as 'effect end' absent from most ordinary cases of perceiving.

2.3 Is the effect view an empirical truth?

If the truth of the effect view – and so the falsity of the simple metaphysical picture – cannot be established on conceptual grounds, we might nonetheless look for empirical evidence in favour of the effect view over the simple metaphysical picture. Given how the effect view and the simple metaphysical picture differ – namely, in the relation between experience and the worldly process involved in perception – the proponent of the effect view ought therefore to look for evidence that experiences are distinct from and caused by their objects. In this regard, we can envisage broadly three kinds of evidence that might be sought by the proponent of the effect view:

- 1. Evidence that experiences occur later than their objects (the temporal stage of the object seen or, where the object is a momentary event, the event itself).
- 2. Evidence that a given kind of experience (as of a red tomato perhaps) can occur (or be made to occur) despite the absence of an appropriate physical relation between the subject and an appropriate object (a red tomato).
- Evidence that the phenomenology of experience is better correlated with, and explained by, the nature of the physical subject (by the workings of her perceptual systems) than by the physical properties of the experience's worldly objects.

The first and second kinds might be considered *direct* evidence for the role of experience as an effect of the object, inasmuch as we are treating the experience the same way as we would any physical object whose role as an effect in some process we were trying to establish: we are looking for the tell-tale signs of temporal separation between

cause and effect, and for the obtaining of an appropriate counterfactual relation between them.

To find the first kind of evidence – evidence for a certain temporal relation between experience and its object – we would first need to establish a way of settling the temporal location of experience. We must of course distinguish between the timing of the event seen and the timing of the seeing of the event. Establishing the temporal location of the object is no problem, at least in principle: armed with appropriate measuring equipment and knowledge of the speed of sound and light, etc., we should have little difficulty working out the temporal stage of the object seen or the time at which the seen event occurred.

We can similarly narrow down the time at which the object is seen – the temporal location of the *seeing* rather than the seen. Doing so requires that we specify what we mean by 'seeing' (or 'perceiving' more generally). In setting out the natural view and the simple metaphysical picture that accompanies it, I claimed that what we call a subject's 'perceptual experience' of an object is nothing more than the state of affairs that is the subject's perceiving the object or the object's being perceived by the subject, a state of affairs describable in physical terms as a causal process relating object and subject in the appropriate way. If that is how we define an episode of perceiving then it is not obvious that the perceiving (the perceptual experience) can be pinpointed to a precise moment in time. It is after all describable as a process which is itself temporally extended. Nonetheless, we do ordinarily talk of seeing objects at a certain time, where the relevant time is that of the successful achievement of the subject in visually apprehending the objects. 'To see', after all, is a success verb and so denotes an achievement of the seeing subject, something the subject *does*. Clearly, the subject cannot do anything purposeful in respect of the object seen until and unless the object affects her, so we might say that it is the time of this physical effect on the subject that is – roughly – the time of the seeing in this sense. (I say 'roughly' because it might be debated which of various physical effects within the subject's visual system must occur for us to grant that seeing has occurred – not the mere stimulation of the retina, certainly, but there may be disputes about the extent to which certain physical upshots involved in recognition must have occurred in order for something truly to have been seen.)

Thus, although in one sense a subject's perceiving an object is a temporally-extended state of affairs involving the object itself, there is a time at which the state of affairs (*qua* process) *as a whole* comes into being, and that is with the occurrence of whichever later part of the process makes the whole process describable as one of perception. Since the causal process between object and subject takes time (negligibly in the case of seeing very close objects but hugely in the case of seeing distant celestial object) the seeing and the seen occupy different temporal locations in all cases. Consider viewing the Andromeda galaxy through a telescope – you see the galaxy *now*, but what you are seeing is the galaxy as it was two-and-a-half million years ago. Indeed, certain celestial objects may no longer even *exist* by the time they are seen.

Now, if we think – as per the effect view – that the seeing here involves the having of an *experience* in some substantial sense, i.e. some state or event within or otherwise 'belonging to' the subject, then it would seem obvious that a supernova which ceased to exist thousands of years before I was born cannot be a constituent of my current experience. However, once we reject the effect view and its substantial conception of experience, there remains no problem about how we can see some distant and now non-existent object – it is simply a matter of our being appropriately causally related to that object.¹⁷

How we respond to the supposed problem arising from perceptual time delay thus depends entirely on our prior theory of perception and the nature of 'experience'. If we think, perhaps on other grounds, that experience is realised by or identical to some state of the brain then we will find further support for the effect view; if, however, we favour the simple metaphysical picture then the (direct) perception of distant and even no-longer-existent objects proves no more problematic in principle than perception of things currently in front of us. Theory thus drives our interpretation of the physical evidence, rather than *vice versa*. (Direct perception of the past has implications for our metaphysical picture of the world's temporal structure, implications I will consider in chapter 7.)

¹⁷ For the same reason, we needn't be put off the natural view by any suggestion that one's experience, in the case of seeing the Andromeda galaxy, must therefore 'reach back' into the past. This way of talking – reminiscent of Campbell's claim that "experience...reaches all the way to the objects themselves" (2002b: 136) – is misleadingly metaphorical: nothing is 'reaching back' in any literal sense, and perceptual experience is not *a something* that stretches backwards in time to embrace its object.

The first line of possible empirical evidence is therefore unpromising. What about the second? Is there evidence that a certain kind of experience can be had in the absence of an appropriate causal link between the subject and an appropriate worldly object? Thus, it is often claimed that whatever brain state a hallucinating subject is in must be sufficient for the occurrence of the relevant experience since, after all, the experience is taken to lack a worldly object. If that is so, and we assume that the subject's brain might be in a similar, indeed exactly similar, brain state when she perceives, then it would seem that we should extend the sufficiency claim to perceptual experience too (see e.g. Sollberger 2008: 4). The argument assumes that, since the same brain state occurs in the relevant case of genuine perception, it must be sufficient for the experience in that case too, such that appeals to the constitutive role of worldly objects are not required to explain the occurrence of the experience. The conclusion of this argument is that hallucinations and perceptions that share a brain state type must therefore share what we might call the experiential element – the experience itself with its phenomenology.¹⁸

Lowe offers what he calls a 'cut-off argument' for the same conclusion. Here, the subject is allowed to look at some object, and has her brain state held fixed while the object is removed (Lowe 2008: 97). The causal connection between the object and subject is thereby cut off, but it is taken that the subject continues to enjoy an experience as of the object. Lowe presents this as providing intuitive support for a causal theory of perception rather than a deductive argument for the sufficiency of a brain state for the having of an experience. Indeed the sufficiency claim is assumed rather than demonstrated. All that the opponent of the causal theory need do here, to ward of the arguments of Sollberger and Lowe, is provide reasons to doubt that a brain state should be considered sufficient for a hallucinatory experience.

My own explanation of hallucination, using the simple metaphysical picture, will seek to do just that. I will show in chapter 4 that we can explain hallucination consistently with the natural view as awareness of worldly objects and their qualities, albeit objects that are not *currently* present before the subject. In the meantime, therefore, I will defer consideration of this second line of evidence, with a promise to tackle it later.

¹⁸ This appeal to sameness of brain states between hallucinations and perceptions plays an important role in Martin's defence of naïve realism, as we shall see later (section 4.5).

Let us turn, then, to the third kind of possible evidence for the effect view, namely evidence for a correlation between the phenomenology of experience and internal states of the subject. Of course, we would expect such a correlation to obtain even on the natural view. Given that our senses enable us to respond purposefully to a wide range of environmental stimuli, we should not be surprised that there is a correlation between the worldly qualities presented to us in perception and what goes on inside our brains when we see those qualities. The distinctive claim that might be thought to favour the effect view is that the phenomenal qualities of experience are *better* correlated with goings-on in the subject's brain than with the scientifically discoverable properties of the worldly objects in which those phenomenal qualities seem to inhere.

Whereas the first and second kinds of evidence would be considered *direct* evidence for the effect view, the third kind of evidence is more indirect. Rather than seeking to demonstrate that there is 'an experience' occupying the effect-end of the perceptual process, the third kind of evidence would consist of a correlation between some properties of experience and internal states or events within the subject that might be considered to constitute or realise that experience. The evidence that is being sought is therefore for an *identity* (or, more weakly, supervenience) between experiences and brain states, or perhaps a causal link.¹⁹

In fact, there *is* considerable and compelling evidence for this sort of special correlation between phenomenology and brain states. For example, it is well known that the nature of perceived colour – captured in terms of the structure of similarity relations among colours – is not correlated with any physical properties of ostensibly coloured surfaces or light. Rather, it seems that this structure of similarity relations among colours reflects (is correlated with) the way in which the human visual system processes visual stimulation by light of different wavelengths. Even more strikingly perhaps, a correlation between perceived ('phenomenal') colour and these inner workings of subjects' physiology is adduced to explain the observation that different normal subjects will often see the same objects as being differently coloured.

¹⁹ One might thus claim that brain states *cause* experiences as the next step in the causal process, although this is at odds with physicalism and raises the puzzle of how something physical could cause something that is not itself physical but 'mental'.

These features of colour experience will be the focus of much discussion in the following chapters (see especially chapter 6 and chapter 7), so I will defer consideration of it until then. However, I will argue there that the evidence for what I will call the *subject-relativity* of certain phenomenal qualities need not be taken as evidence for the effect view or, more generally, for the claim that those phenomenal qualities are themselves internal to the subject in some way. However, resisting the effect view and upholding the natural view in the face of the manifest subject-relativity of colour (not to mention certain other perceptible qualities) will force us to commit to some surprising and far-reaching conclusions about the nature of the world and our place in it as subjects.

CHAPTER 3 - PHENOMENOLOGICAL PROBLEMS (I): ILLUSION

3.1 Introduction: the arguments from illusion and hallucination

The natural view is, above all, a claim about the phenomenology of perception, namely that what we call the phenomenal character of experience is really just the sensory character of the worldly scene perceived. I have presented this claim as a philosophical gloss on a common-sense or pre-philosophical intuition about perception, namely that perception is just a straightforward and immediate confrontation with objects in our environment. Nonetheless, the fact that such a view has been so rarely defended is due to certain experiences whose phenomenology does not appear attributable, in whole or part, to how things are before the subject. In this and subsequent chapters I will describe such experiences and the problems they seem to present for a proponent of the natural view. Notably, these experiences include those that we call *perceptual illusions* and *hallucinations*, although there are aspects of the phenomenology of ostensibly veridical experiences that also present problems for the natural view.

Traditionally, objections to the sort of perceptual immediacy postulated by the natural view have rested on two observations: first, that the phenomenology of perceptual experience sometimes fails in some sense to 'match up with' how things are in the subject's environment; and, second, that a genuinely perceptual (and hence ostensibly world-presenting) experience could have the same phenomenology as an experience in which nothing worldly is presented. The first observation concerns of course the possibility of *illusion*, while the second concerns the possibility of *hallucination*. Note that perceptual illusions are clearly cases of genuine (if not veridical) perceptual experience, since these are cases in which an object is perceived, albeit misleadingly. Hallucinations, by contrast, are by definition cases in which there is the appearance of an object despite there being no corresponding object that is perceived.

These then are the two broad classes of experience that are widely held to undermine naïve realism:

<u>Perceptual illusion</u>: the phenomenal qualities manifest in experience sometimes fail to match up with or correspond to the way the perceived objects really are.

<u>Hallucination</u>: some experiences seem to present phenomenal qualities even although those experiences lack a worldly object; in such cases the hallucinatory phenomenology cannot be attributed to the objects 'out there' in the subject's environment.

Experiences belonging to one or other class become problematic for the natural view once we use them as the basis of a certain kind of argument, one which draws a conclusion about the nature of specifically illusory or hallucinatory experience and applies it to ordinary, veridical perceptual experience. This kind of argument comes in two versions, often called the 'argument from illusion' and the 'argument from hallucination' respectively. These arguments have the same overall structure, allowing them to be expressed in amalgamated form as follows:

The Argument from Illusion/Hallucination

- 1. For every veridical perceptual experience there could be an illusory or hallucinatory experience that seems the same to the subject.
- 2. The best explanation for the fact that two experiences seem the same is that they have the same metaphysical nature.
- All illusory and hallucinatory experiences involve awareness of phenomenal qualities that are not instantiations in worldly objects; those qualities are not 'out there', inhering in worldly objects.

Therefore

4. All veridical perceptual experiences involve awareness of phenomenal qualities that are not instantiations in worldly objects; those qualities are not 'out there', inhering in worldly objects.

Premise 1 derives from a plausible claim about the possibility of illusory or hallucinatory experiences that subjectively 'match' any veridical perceptual experience. A subjectively matching experience we can take to be one that the subject could not tell apart from a veridical perception by reflection alone, even assuming she had ideal powers of reflection on her own experiences.²⁰ Premise 2 is an abductive inference from a certain phenomenological equivalence of experiences (their subjective indiscriminability) to their metaphysical equivalence. That is, given a triad of matching experiences – one a veridical perception, one an illusion and one a hallucination – we might expect a single explanation to account for their common appearance. From here, one only needs to assume that illusory and hallucinatory experiences involve awareness of phenomenal qualities that are not 'out there' in the environment (premise 3) to derive the conclusion that they must not be 'out there' in the veridical perceptual case either.

The overall thrust of the argument is towards a generalising or unifying explanation for the similarities among diverse experiences, a unifying strategy that is made explicit in the abductive inference that forms premise 2. We might find support for this unifying strategy in the fact that some non-veridical experiences may be only partial, so that we perceive an object's shape accurately but not its colour (Smith 2002: 44), or hallucinate an object against a veridically perceived background. That veridical and illusory qualities can seemingly coexist within a single experience might be taken to reinforce the idea that they have the same metaphysical nature. Just as compelling in this regard is the possibility of experiencing a smooth transition from a veridical perceptual experience to an illusory experience (ibid.). In his sense-datum version of the argument from illusion, Robinson makes the generalising step as follows: "There is such continuity between those cases in which objects appear other than they actually are and cases of veridical perception that the same analysis of perception must apply to both" (Robinson 1994: 58).

Having thus denied the natural view, the proponent of the arguments from illusion and hallucination must offer an alternative explanation of the phenomenology that is shared by matching experiences of each kind, whether veridical, illusory or hallucinatory. One approach has been to suggest that the phenomenal qualities manifest in illusory and/or hallucinatory experiences, not being 'out there', must instead be 'in here', qualifying some mental or mind-dependent objects ('ideas' or 'sense data'), or qualifying experience itself (i.e. 'qualia').²¹ Since the awareness of such mind-dependent objects or qualities suffices,

²⁰ This appeal to indiscriminability by reflection alone is a central feature of Martin's (2004) account of hallucination, as we shall see in chapter 4.5.

²¹ For a recent defence of sense-datum theory, see Robinson (1994). For defences of qualia as intrinsic nonrepresentational properties of experience, see e.g. Block (2003) and Shoemaker (1991).

in the case of illusion or hallucination, for the subject's having an experience with a certain phenomenology, it is inferred (in the pursuit of a unified explanation) that their occurrence might also explain the phenomenology of the matching veridical experience. On this view, the phenomenology of a perceptual experience must be constituted not by worldly objects and their qualities but by some mind-dependent objects or qualities. What makes the experience genuinely perceptual is then likely to be attributed to the obtaining of an appropriate causal relationship between the occurrence of the sense data or qualia and the worldly object perceived.

An alternative way to accept the conclusion of the arguments from illusion or hallucination rejects the notion that we are aware of inner objects of awareness (sensedata) or intrinsic qualities of experience (qualia). On this alternative view, known as representationalism, we are aware of worldly objects by dint of representing the instantiation of certain qualities. What's more, the representationalist might further argue that when we represent the instantiation of these qualities it is the qualities themselves that constitute the phenomenology of our experiences (Dretske 1995, Lycan 2001, Byrne & Tye 2006). Our awareness is thus taken to be 'direct' in the limited sense that it is not mediated by awareness of anything else (sense-data, qualia, etc.). Crucially, however, the representationalist can allow that a subjectively matching triad of veridical perception, illusion and hallucination involves experiences that have the same metaphysical nature. In each case the experience represents the instantiation of certain qualities, and it is this representation that determines the experience's phenomenology. The difference between the veridical and the non-veridical experiences is just that only in the former case are all the represented qualities are *actually* instantiated as they appear (Fish 2010: 66, 79). It remains that the actual instantiation or otherwise of the properties is extrinsic to the experiences' being as they are, so a common metaphysical explanation is available for their mutual indiscriminability.

Although the arguments from illusion and hallucination turn on the desirability of a unifying explanation for the phenomenology of diverse experiences – some veridical and some non-veridical, and some of the latter not even genuinely perceptual – it is questionable whether the arguments' proponents can truly be said to offer a genuine explanation for even the non-veridical or non-perceptual cases. As Putnam points out with respect to sense datum theories, the purported explanations appeal to entities or processes –

sense data or representations – that are either quite mysterious or difficult to square with our ordinary and indeed scientific picture of the world and our place in it (Putnam 1999: 29). To that extent, then, the pivotal step in the arguments relies more on the negative assumption that the phenomenology of non-veridical and/or non-perceptual experiences *cannot* be explained as relations of awareness holding between subjects and worldly objects. It is this negative assumption that I will seek to overturn, with the help of the simple metaphysical picture. That is, I will argue, rather counterintuitively, that illusory and even hallucinatory phenomenology can be explained wholly in terms of relations of awareness holding between subjects.

To that end I will, in this chapter and the one that follows, assess the prospects for reconciling the natural view with illusion and hallucination, respectively. As we shall see, one way naïve realists have accounted for illusion is by rejecting premise 3 of the argument from illusion, namely the claim that illusory phenomenal qualities must not be 'out there' inhering in worldly objects. According to this counter-argument, the illusory aspect of the experience resides not in the sensory qualities presented in perception but instead in some associated cognitive functioning, including our inclination to form judgments based on what is presented. Favourably for the natural view, this allows us to defend the claim that, even in perceptual illusion, the phenomenology is constituted by worldly objects and their qualities. Although I will indeed endorse this counter-argument as a way to preserve the natural view in the face of *certain kinds* of perceptual illusion, I will note that certain other kinds of illusion may not be so easily accommodated. Certain optical illusions, for example, seem to involve genuinely phenomenal features of experience that cannot be explained without appealing in some measure to what goes in within the subject's perceptual system. These therefore threaten what we might call a 'retreat inwards' - an abandonment of the natural view's commitment to explaining the phenomenology of experience by reference to worldly objects and qualities. The reason that this threatens abandonment of the natural view and not just its supplementation is the thought that if goings-on internal to the subject can be shown to be sufficient for some aspects of phenomenology then we might as well treat it as sufficient for all of it.

Having acknowledged this residual problem posed by illusion, I will then turn (in chapter 4) to the even more obviously troublesome phenomenon of hallucination. On the face of it, hallucination presents the starkest challenge to the natural view, since here we

seem to have a class of experiences that possess phenomenology while lacking the worldly objects and qualities that would, in the perceptual case, account for that character. The standard naïve realist response to the argument from hallucination is to reject premise 2 of the revised argument, i.e. the claim that subjectively matching perceptual and hallucinatory experiences are best explained by appeal to a shared metaphysical nature. This leads the naïve realist to one or other form of disjunctivism about perceptual experience, as I will explain. Disjunctivist accounts therefore eschew a unified explanation of ostensibly similar experiences, a strategy that brings some serious shortcomings. Here, I will argue that the simple metaphysical picture points to an account of hallucination that avoids these shortcomings and, more positively, gives a unified account of perception and hallucination. The strategy is to reject premise 3 of the revised argument while endorsing premise 2 -that is, to argue that veridical perceptions and hallucinations have the same metaphysical nature insofar as both are object-involving. This of course is the precise opposite of the standard disjunctivist defence of naïve realism. Defending the natural view by claiming that the phenomenology of hallucinations is constituted by worldly objects and their properties might seem a thoroughly perverse strategy. My task will be to show how the simple metaphysical picture can be used to make the strategy intelligible and even plausible.

It is important to acknowledge that, besides illusion and hallucination, there are certain phenomenal features of *veridical* perceptual experience that seem to challenge the natural view. In chapter 5 I will show how the simple metaphysical picture-based explanation of hallucination points to an explanation of what is plausibly a ubiquitous feature of veridical perception, namely the influence of past perception on the phenomenology of current perceptual experience via perceptual learning and belief. In particular, the role of perceptual learning and beliefs in determining the phenomenology of perceptual experience – including the phenomenon known as *cognitive penetration* – is increasingly recognised, and any philosophical account of perceptual experience will be the stronger for explaining it. As I noted in chapter 1, doing so compatibly with the natural view might appear especially challenging since it requires us to reconcile two features of perception – its immediacy and its generality – that might seem at first sight to be somewhat in tension. Nonetheless, I will argue that the object-involving account of hallucination paves the way for an explanation of perceptual learning and cognitive penetration that similarly appeals to nothing other than worldly objects and qualities as the bearers of phenomenal character.

Finally, in chapter 6, I will turn to certain features of colour experience that seem to pose further difficulties for the natural view. In particular, I will consider cases in which the colour component of an experience's phenomenology cannot easily be explained as the presentation of qualities inhering in the subject's environment. Such cases suggest that the subject plays a constitutive role in the nature of the phenomenology of experience in a way that conflicts with the supposed explanatory virtue of naïve realism, i.e. that qualities inhering in the environment play a constitutive role in the phenomenology of experience but not *vice versa*. Again, I will suggest that the simple metaphysical picture points to a possible way to accommodate this, although in doing so it radically overturns orthodox assumptions about the metaphysical priority relation holding between processes and their constituent parts. The implications of this will become the subject of chapter 7.

3.2 Illusion and the natural view

It is tempting to begin with illusion, if only because it appears at first sight to present a less stark challenge to the natural view than hallucination. After all, perceptual illusion remains a case of perception, since there is still a worldly object of which one is aware. It therefore remains true that there is something 'out there' to which we might attribute phenomenal character. What illusion does, however, is throw back into question the nature of this correspondence between what is out there and this phenomenal character. According to the natural view the object or scene perceived is what possesses the phenomenal character. However, the troubling feature of illusory experience is that the scene's phenomenal character is in some way misleading about the nature of its bearer. In the argument from illusion as I presented it above, the driving force is the observation that in illusory experience worldly objects' actual qualities and the qualities they *appear* to have that is supposed to refute naïve realism.

More specifically, what is of interest here is whether the misleading character of illusory experience is problematic for the natural view. We might have good reason to think so if the phenomenology of illusory experience is in some cases not attributable to the objects and qualities inhering in the subject's environment. The natural view allows us to explain how things appear solely by reference to the nature of those things themselves, the conditions of perceiving and the subject's perspective. If, for example, we have to

appeal to internal goings-on in the subject to account for some aspects of phenomenology, whose remainder is attributable (as per the natural view) simply to worldly objects and their qualities, then we seem to have two very different factors jointly constituting what appears on reflection to be a seamlessly unified phenomenon. This is where the pressure to generalise arises – to give a unified explanation of an ostensibly unified phenomenon. As such, if illusion proves that phenomenology is at least in part not attributable to worldly objects and their qualities then the defender of the natural view has more work to do to resist that urge to unified explanation.

We might therefore envisage two broad strategies for explaining perceptual illusion consistently with the natural view. The first would keep phenomenology and illusion apart, explaining the former exhaustively in terms of how things are perceptibly before the subject, and explaining the latter by appeal to non-phenomenal factors. The second strategy would acknowledge that the illusory aspect of experience makes a real difference to phenomenology, but explain it in a way that does not threaten the natural view. The second strategy looks unpromising at first sight, since any putative illusory aspects of phenomenology are assumed to be generated internally, an assumption that invites one to generalise to an internalist explanation for all phenomenology, illusory or otherwise. Consequently, naïve realists tend to favour the first strategy, as we shall see.

To some extent, whether we locate the illusory aspect of experience within its phenomenology depends on what kind of illusion is involved. The category of perceptual illusion includes a diverse range of experiences, not all of which operate via the same mechanism. In some the illusory aspect is less to do with how things appear than with the way we *respond* to how they appear – these cases are amenable to the first explanatory strategy described above. In other cases it is harder to deny that the illusion is a genuinely phenomenal matter, and here some form of the second strategy will be required. In the following sections I will make use of a simple taxonomy of perceptual illusion offered by Fish (2009: ch. 6), and show how different kinds of illusion demand different kinds of explanation.

3.3 Physical illusions

Fish calls 'physical illusions' those that can be explained purely by reference to the object perceived, its perceptible qualities and its location relative to the subject (2009: 150 ff.). These are cases, in other words, that are consistent with what I have called naïve realism's explanatory virtue. To cite a classic example of visual illusion, a straight stick partly submerged in water looks crooked. Experience is thus potentially misleading, if it leads us to think that the stick is, in itself, crooked. Nonetheless, the fact that the stick appears crooked under these circumstances is perfectly explicable in terms of the relative refractive indices of air and water, coupled with the spatial relation of the stick and water to the perceiving subject. Everything that we need to explain the appearance of crookedness is 'out there' in the environment, compatibly with the natural view (cf. Austin 1962: 26). Or consider the case of seeing a white cube illuminated by red light. It might appear as though the cube is actually red, in which case one's experience has proved misleading. All the same, we can perfectly well explain how this misleading appearance has come about, purely by reference to the qualities of the incident light and the cube's surface.

The argument from illusion could exploit such physical illusions only if the natural view insisted that things must always appear to be the way they actually are – that straight sticks must always appear straight, red things always red, and so on. But the natural view makes a weaker demand: just that the overall phenomenal character is borne by the worldly things themselves. This leaves room for the phenomenal character of our environment to mislead us in some way. As such, one appealing way to defend the natural view against the argument from illusion is to reject as illegitimate the inference that, because phenomenal character can sometimes mislead us about the true nature of worldly things, that phenomenal character must not be 'out there' in the way the natural view requires. If phenomenal character sometimes misleads, then it misleads by presenting the worldly objects as possessing some qualities they really lack. This was why, in section 1.5, I expressed doubts about my first pass definition of the natural view as the claim that phenomenal qualities are 'out there', inhering in the worldly objects perceived.

In physical illusion, then, the illusion properly resides in some misjudgment about how individual elements of a scene must be in order for the overall (phenomenal) character to be that way. This has led some to suggest that there is a sense in which the

phenomenology, in the examples cited, is not illusory at all – rather, it perfectly well presents the scene as it is (see e.g. Travis 2004). *That* is just how a straight-stick-in-water should look (somewhat like a crooked stick), and *that* is how a white-cube-in-red-light should look (much like a red cube in white light). The phenomenology itself is not illusory – for that, *ex hypothesi*, is just a matter of how the scene is within the subject's perspective. That being so, no unifying explanation of phenomenology is required because what is illusory in certain experiences is not to be found in their phenomenology at all.

But where else could it be found? I suggested above that the illusion results from their tendency to *mislead* us. Specifically, we are misled into making erroneous judgments about how things must be for things to appear as they do. One plausible option is thus to say that the illusion properly resides in the thoughts that we are inspired to form on the basis of what is presented to us perceptually. This view recognises that it is perfectly possible that an experience may accurately exhibit the phenomenal character of a subject's perceptible surroundings and yet, at the same time, should mislead us about the true nature of the objects and qualities occupying those surroundings. For example, we might say the appearance of a stick in water is illusory if it inclines us to think that the stick really is crooked, or that the appearance of the penny is illusory if, seeing it from an oblique angle, we take it to be elliptical rather than circular. (Of course, we see circular things from oblique angles all the time and rarely misjudge them to be anything other than circular, testament to the power of our cognitive faculties to lead us in the right direction more often than not.)

The claim that phenomenology is not inherently illusory is defended by naïve realists as a corollary of their more general insistence that perceptual experience is not essentially representational. Phenomenal character is thus nothing more or less than the sensible appearance of worldly objects and qualities and entails no commitment to how things must be for them to appear the way they do. Travis quotes Austin approvingly in making this point: "our senses are dumb [...they] do not *tell* us anything, true or false" (Austin 1962: 11; quoted in Travis 2004: 64). Perception "merely bring[s] our surroundings into view", leaving us free to judge how things are in those surroundings; the experience itself does not commit us to a certain judgment about how things are, but simply presents the things themselves and thereby *makes it possible* for us to judge one way or the other (Travis 2004: 64).

What emerges is what I will call a *two-faculty view of illusion*, namely the view that illusion results from the interplay of these two more or less distinct faculties:

- 1. A faculty of awareness which acquaints the subject directly with the worldly objects and their qualities.
- 2. A cognitive faculty that inclines the subject to judge (mistakenly) that the objects or their qualities are some way.

Brewer thus asserts that there are "two levels in the subjective character of experience" (2006: 172). The first level is the presentation of the mind-independent object that plays a constitutive role in the phenomenology of the experience. This yields what Brewer calls the 'core subjective character' of the experience. The second level is a cognitive or conceptual one in which the mind-independent object is recognised or categorised by virtue of its "visually relevant similarities with paradigms of various kinds of such things" (Brewer 2007: 91). Brewer states that such paradigms are "instances of the kinds in question, whose association with the terms for those kinds partially constitutes our understanding of them, given our training in the acquisition of the mind-independent object from some contingently connected tendency to judge which state of affairs obtains given the phenomenology of the experience constituted by the presentation of that object. He allows that there is mental representation, albeit that representation or judgment is a cognitive act distinct from the perceptual experience itself (Travis 2007: 233 ff.).

It is important to note that, to give rise to illusion, the misleading object of awareness need not actually issue in a false judgment. Rather, it need only *incline* the subject towards such a judgment. Thus, Brewer's account requires only that the object, presented a certain way, "has the evident power to mislead *her*, whether or not this error is actualized in any false judgment" (2006: 169). This allows us to acknowledge that some illusions retain their illusory appearance even when we know that it *is* an illusion; the workings of our cognitive faculty may *incline* us to believe that one line of the Müller-Lyer illusion is longer than the other even though, as seasoned philosophers of perception, we know better than to acquiesce in the judgment that it is so.

Crucially, the two-faculty view allows the naïve realist to separate the illusory aspect of experience from the mere awareness of worldly particulars and their qualities, i.e. to hold that the illusion is properly attributable to flaws in the act of recognition or judgment rather than the awareness on which that recognition or judgment is based. It therefore accords with the first strategy mentioned above by which we might defend the natural view from the argument from illusion. As such, it fits well cases of physical illusion, in which the phenomenology is wholly explicable in terms of the worldly objects perceived, their qualities and their spatial arrangement in respect of the subject. Here, the subject's surroundings are simply presented to the subject perceptually and it is in virtue of this perceptual presentation that the constituents of those surroundings are available as the objects of recognition and judgment.

As we shall see, this 'divide and conquer' strategy, in appealing to two kinds or 'levels' of mental state to account for the illusory character of experience, is key to naïve realists' explanation of other illusions besides physical illusions. There is a potentially problematic difference, however, in how the two-level approach accounts for the different kinds of illusion. In physical illusions the two levels – the mere awareness and the tendency to judge – seem readily distinguishable: the phenomenology is readily explicable in terms of the layout of the scene perceived or the conditions of perceiving, while any 'illusion' can be understood as a subsequent misjudgment based on that phenomenology. In certain other kinds of illusion the distinction appears less sharp: in these cases there appears to be a role for the (ostensibly inner) cognitive states in actually influencing the phenomenology of the illusory experience, a role which threatens the natural view by resurrecting the unifying explanatory strategy that is key to the argument from illusion.

3.4 Cognitive illusions

On the face of it, the two-faculty view seems tailor-made to account for Fish's second class of perceptual illusions – those he calls *cognitive illusions* (2009: 165 ff.). These are cases in which the illusion arises because of miscategorisation, in other words the erroneous application of what I have called the subject's cognitive faculty and Fish more precisely calls a subject's *conceptual-recognitional capacity*. A conceptual-recognitional capacity is taken to be the subject's ability to recognise some *F* as an *F*, and cognitive illusion results from deploying a conceptual-recognitional capacity for *F*s when perceiving

a non-*F*. To use his example, a subject is prey to a cognitive illusion when he sees a coil of rope and mistakenly judges it to be a snake. In other words, the subject sees the rope but sees it *as* a snake. As Fish points out, in order to misperceive the rope in this way, the subject must possess the concept of a snake, hence the term 'conceptual-recognitional capacity' (2009: 167). It is some degree of similarity between the perceptible qualities of the rope and those of paradigmatic snakes that induces the subject to deploy the inappropriate conceptual-recognitional capacity (perhaps encouraged, as Fish suggests, by some priming – a prior warning to look out for snakes, say). The resemblance to Brewer's account is obvious: Brewer talks of illusion as arising from the subject's recognition of an object's visually relevant similarities with paradigms of some other kind of object (Brewer 2007: 91).

Now, insofar as a rope might look as though it were a snake, we might seem to be imputing a 'snake-ish' element in the experience's phenomenology. If there were indeed such an element, it would seem that it must be explained by the two faculty view as the product of the cognitive faculty and not the faculty of direct awareness. A role for the cognitive faculty in determining the phenomenology of illusion is admitted by Brewer: although he claims that direct awareness of the object delivers the experience's 'core' phenomenology, he does allow that a subject's cognitive engagement with the object can have a bearing on how that object looks, and therefore on the experience's (wider) phenomenology. He says that "[t]his is the phenomenology of conceptual categorization, or recognition, not that of basic experiential presentation...[but] it is aptly titled *phenomenology*, all the same" (Brewer 2007: 93).

Brewer's admission that the cognitive component of the illusory experience contributes some phenomenology leaves us with both a puzzle and a worry. The puzzle is how the cognitive phenomenology relates to that yielded by mere awareness. If they are different kinds of phenomenology, how do they interact? Does the former somehow alter or skew the latter, or does it positively add some *extra* phenomenology (some extra colour perhaps)? Either way, we need some explanation of how an ostensibly internal faculty of cognition can interact with the core subjective character consisting just of the worldly objects perceived. If this expresses the puzzle then the worry, as noted earlier, is that an account which allows internal processes to positively contribute to phenomenology threatens to resurrect the argument from illusion and its unifying explanatory principle. In other words, if internal goings-on can generate *some* phenomenology then why not all of it?

Fish likewise holds that the application of a conceptual-recognitional capacity influences the phenomenology of experience, but argues that it does so by influencing what in the world a subject perceives. On this view, the cognitive faculty is not itself a determinant of phenomenology, but makes available certain worldly features which do contribute to a scene's phenomenal character. He explains the role of the conceptualrecognitional capacity in terms of perceiving *facts*, where by 'fact' he means a structured complex of objects and properties (i.e. that which is sometimes called a state of affairs) (Fish 2009: 52-3). This, for Fish, captures our most basic way of perceiving the world, namely as objects bearing qualities.²² To see a rope as a rope is, on this view, to recognise the obtaining of the fact that is a rope's being present. Succeeding in this way to perceive the fact of the rope's being present results when the various qualities of the rope appropriately trigger the subject's conceptual-recognitional capacity for ropes. Meanwhile, when the subject falls prey to a cognitive illusion and misperceives the rope as a snake, she clearly does not perceive any snake-ish fact, for there is none to be perceived; rather, she erroneously applies her conceptual-recognitional capacity for snakes instead (Fish 2009: 168).

This shift in the application of conceptual-recognitional capacities is supposed, on Fish's view, to account for what we might loosely call a phenomenological difference between the illusory and veridical experiences. As he explains it, all the difference is accounted for by the subject's being aware of something extra in the veridical case, namely the fact of the object's being a rope. Perceiving this fact imparts to the experience the "further phenomenal property" of acquainting one with that fact (2009: 167). This does not, however, amount to the claim that there is a change to the qualities of which the subject is aware. We need to be aware here of Fish's distinction between what he calls *presentational character*, which is constituted by the perceived scene itself (and so is close to my 'phenomenal character') and what *he* calls phenomenal character, namely "the

²² In stating that facts, so defined, are the "basic constituents of presentational character" Fish emphasises that the claim is phenomenological rather than ontological; that is, it is a claim about *how things seem to us* when we perceive – first and foremost as whole states of affairs – and makes no claims about the ontological order of priority between facts and their constituent objects and properties (Fish 2009: 53). He nonetheless insists that facts are "pieces of reality" that can be perceived, a claim that causes problems for his account, as I will explain.

property [of an experience] of acquainting the subject with such-and-such a presentational character" (2009: 15). Thus, when he talks of experience "acquir[ing] a further phenomenal property" he does not imply that there are any additional phenomenal qualities (in my sense) of which the subject is aware.

We might thus find in Fish's fact-based account a way to understand Brewer's claim about cognitive 'phenomenology'. If we think of phenomenology in terms of the way things appear or (in the visual case) look to the subject, then we can understand Fish as claiming that the rope looks different to the person who sees it as a rope from how it looks to the person who sees it as a snake, not because either is aware of any different objects or qualities, but because one recognises the obtaining of the relevant fact and the other does not.²³ This point about what one is aware of is crucial: in neither the illusory nor the veridical perception does the cognitive faculty add any phenomenal qualities to the experience; one is aware of just the same qualities (and objects) in each case. Instead, the cognitive faculty contributes to the experience by enabling us to see how the various phenomenally presented elements of a scene – the various perceptible objects with their qualities - fit together as constituents of one or more facts. In the veridical case the subject correctly recognises them as constituting the actually obtaining facts, such that she becomes aware of those facts as facts. In the illusory case, meanwhile, the subject erroneously applies some other conceptual-recognitional capacity and so comes to misperceive the presented objects and qualities as constituting some non-obtaining fact(s). Fish's account thus succeeds in attributing the illusory aspect of the experience to the (incorrect) application of a certain cognitive faculty, and not to any change in what objects or qualities the subject is aware (in other words, without altering what Brewer would call the 'core subjective character').

There are some difficulties with Fish's account, however. For one thing, it seems to deny any distinctive phenomenology to illusory, as opposed to veridical perceptual experiences. I said above that all of the phenomenological difference between the veridical and illusory experiences is, on Fish's view, accounted for by the extra awareness in the veridical case of the relevant fact. This is because there is no distinctive fact of which the subject can be aware in the illusory case – the object is a rope and not a snake, so the

²³ As I will explain below, Fish sees the same mechanism at play in veridical perception, arguing that it shows up in the difference between the experiences of an expert subject and a naïve subject when confronted perceptually with some object (Fish 2009: 68).

subject's conceptual-recognitional capacity for snakes fails to latch onto any actual fact. There is therefore nothing extra in the illusory case for the subject to be aware of – no fact of the object's being a snake. This misperception, or deployment of the wrong conceptualrecognitional capacity, does not in itself have any phenomenological implications.

We can see why a naïve realist like Fish should wish to promote this view and to attribute all cognitive phenomenology to the apprehension of genuine worldly objects, qualities or states of affairs (facts). The alternative would be to grant that ostensibly internal processes of cognition are sufficient for some elements of phenomenology, an admission that would threaten to revive the argument from illusion and the unifying strategy of attributing *all* phenomenology to internal goings-on. All the same, a refusal to allow that the illusory aspect of experience brings any distinctive phenomenology risks looking counterintuitive. Consider that Fish uses his fact-based account to explain the plausible phenomenological difference between the perceptual experiences of two subjects viewing the same object under the same conditions: one subject is an expert who recognises the object as being of a certain kind, while the other is a novice who does not recognise the object; Fish uses the example of a scientist and a child looking at a cathode ray tube (2009: 68). The difference between the two experiences is due to the expert's possessing a conceptual-recognitional capacity for the object in question, such that she is able to be aware of or recognise a fact (the object's being a cathode ray tube) that is inaccessible to the novice. Now imagine that we have two subjects looking at some object - a rope, say. The first subject is a real novice - a baby perhaps - and fails to recognise the rope as a rope but is nonetheless able to attend to its visible qualities. The second subject has a peculiar fascination for snakes and is always on the lookout for them; on seeing the rope she mistakenly judges it to be a snake. She thus erroneously deploys a conceptualrecognitional capacity for snakes, and thereby fails to acquaint herself with any facts, in Fish's sense. The first subject, meanwhile, has signally failed to deploy any conceptualrecognitional capacities at all, but has at least avoided illusion as a result. Since both are aware of the same object, and neither is aware of a pertinent fact that is hidden from the other, it seems we have no grounds for supposing that the two subjects' experiences differ in any phenomenological respect. But this seems implausible. Or, at least, it posits an intuitively unconvincing asymmetry between two pairs of experiences: expert perception and novice perception on the one hand; and, on the other, illusory perception and novice perception. If we think there is some phenomenological difference between the
experiences enjoyed by novice and expert perceivers, does it not seem as though there should be some difference between those enjoyed by the novice perceiver and the perceiver subject to illusion? Fish's fact-based account generates the asymmetry, but we lack any independent reason for supposing the asymmetry to hold, and indeed our intuitions might not support it. Surely things just *look* different to the deceived snake-fancier in virtue of (erroneously) deploying her conceptual-recognitional capacity for snakes.

The asymmetry is, in any case, an artefact of a more fundamental problem in Fish's account, one which relates to an ambiguity in his claim that expert subjects, and experts only, can perceive certain facts. Specifically, the ambiguity is in his use of 'perceive' as applied to facts. According to Fish, for a subject to perceive a fact she must deploy the appropriate conceptual-recognitional capacity when confronted by relevant object and qualities. She needn't *know* that the fact obtains; deploying the appropriate conceptual-recognitional capacity is sufficient (2009: 54). The fact thereby comes to constitute part of what Fish calls the presentational character of the experience. The appeal to facts is, as noted above, intended to capture how things seem when we experience: we perceive the world as a collection of facts, as "things bearing properties" (ibid.: 53). However, given Fish's stipulation that facts in this sense are not to be understood as true propositions but as "pieces of reality" (ibid.), we might wonder whether he is entitled to claim that perceiving a fact requires that one possess the relevant conceptual-recognitional capacity. Surely if facts are pieces of reality then they can be perceived irrespective of whether they are thereby *recognised*.

On the face of it, Fish's claim that the successful expert perceiver is acquainted with something in the world – a fact – that is inaccessible to the novice or deceived subject looks like an attractive way to account for the intuitively appealing notion that the same scene might look different to each of these subjects. Importantly for the proponent of the natural view, the difference would be explained by a difference in what in the world is seen. However, once we acknowledge that the difference between the subjects is not in what is *seen* but what is *recognised*, we are left instead in the position we started with, namely that the difference is somehow explained by a cognitive faculty that is presumably internal to the subject. This means that Fish's appeal to perception of facts amounts to little more than a restatement of the claim that our cognitive faculty plays some role in

shaping the way things appear. And this doesn't bring us any closer to an understanding of *how* this cognitive 'shaping' of experience comes about, especially if we are inclined to think that *how things look* is simply a matter of *what things we see* (and the physical and perspectival circumstances in which we see them).

We have, in fact, good empirical reasons to think that our cognitive faculties *can* have a significant determinative influence on how things appear, even to the extent of altering the apparent (phenomenal) qualities of the worldly objects of awareness. In other words, there are good grounds for thinking that the ostensibly 'inner' conceptual-recognitional capacity can influence what Brewer would call the core subjective character of the experience. This is apparent in cases of perceptual learning as well as what is called *cognitive penetration* of perceptual experience, where it seems that our past experiences and beliefs can help to shape the phenomenology of perceptual experience. That being so, the defender of the natural view is again challenged to explain, first, how the faculties of cognition and awareness interact to yield a unified experience, and, second, why admitting a role for ostensibly inner states (of memory and cognition) in shaping phenomenology does not resurrect the argument from illusion.

I will argue later (chapter 5) that the simple metaphysical picture offers an explanation of perceptual learning and cognitive penetration, and one which avoids the argument from illusion by claiming that the elements of phenomenology that they deliver are not internal after all. The explanation of how this might be so relies heavily on the account of hallucination which I will offer in chapter 4, so I will defer it until later.

As it is, there are other, more familiar forms of illusory perception that also put pressure on the claim that cognitive phenomenology can be distinguished sharply from the core subjective character delivered by direct awareness. These include what Fish call *optical illusions*.

3.5 Optical illusions

Optical illusions comprise another class of illusion in which, like cognitive illusions, the inappropriate application of a conceptual-recognitional capacity is supposed to play the decisive role (Fish 2009: 172-181). Fish notes that optical illusions share with physical

illusions the fact that they are likely to be shared by any normal subject confronted with the same scene, and that they are, in this sense, predictable. Cognitive illusions, by contrast, are dependent on the background knowledge, beliefs and expectations of the subject, such that two subjects confronted by the same scene may not share an illusory experience, nor indeed might the same subject have the same illusory experience on two occasions of viewing the same scene (2009: 172-3). What optical illusions share with cognitive illusions, however, is that they cannot be explained solely by reference to the nature of the scene before the subject – its constituent objects, their qualities, arrangement in the subject's perspective, and so on.

Fish sees the Müller-Lyer illusion as an example of an optical illusion: it has the intersubjective and predictable features of a physical illusion, but the seeming difference in length between the two lines is not explicable in terms solely of what is presented to the subject. Rather, to explain the illusion it seems we need to appeal to *how the lines affect the subject*. As Fish puts it, "[t]he reason optical illusions lead to a nonveridical experience is that the relevant features of the perceived scene function so as to trick or mislead our perceptual mechanisms" (2009: 172). It seems in these cases that our perceptual mechanisms – the workings of our sensory organs and subsequent neural activity – do not merely provide the enabling conditions for awareness but play some role in influencing or determining at least some aspects of the phenomenology of our conscious awareness.

Brewer too cites the Müller-Lyer illusion to illustrate his two-faculty account of illusion. Although the two lines appear to be of different lengths, Brewer insists that the apparent difference is not presented to us as an aspect of the 'core subjective character'. Instead, what are presented to us are just the lines themselves, and the illusion of disparate lengths arises from the operation of some other cognitive faculty that inclines the subject towards a mistaken judgment. The cognitive act in question is a recognition that the two lines presented bear some 'visually relevant similarities' to a possible veridical perception that presents lines of different length – in Brewer's example, a set-up in which the longer line is further from the subject than the shorter line, such that their ends subtend the same angle on the subject's eye (2007: 91). Figure 1 below illustrates one such set-up, in which the lines of the Müller-Lyer illusion are conceived as elements in a three-dimensional array of objects: the upper line (bold) becomes the near edge of a box attached to a wall above

the bottom corner of a room, where the lower line becomes the intersection of the wall and the floor.



Figure 1 – three-dimensional interpretation of the Müller-Lyer illusion

Brewer's two-faculty solution here proposes that, when we view the two parallel lines of the Müller-Lyer illusion, there is nowhere any *actual* difference between the phenomenal qualities of which we are aware and which correspond to the length of the lines themselves. This is of course because the correspondence here is *identity*: the phenomenal qualities just are the lines' lengths, 'out there'. The *appearance* of a difference in length is not constituted by what we are aware of, but by some further cognitive act in which the lines are misinterpreted as belonging to three-dimensional, rightangled objects. This explanation finds support from cross-cultural studies in which subjects living in 'noncarpentered' visual environments lacking abundant right-angled joins between surfaces (chiefly those not living in 'boxy' houses) are less susceptible to the illusion (Kitayama & Cohen 2007: 576).

This cross-cultural difference might incline us to assimilate the Müller-Lyer to a cognitive rather than an optical illusion. However, as Fish notes, what is distinctive about optical illusions as opposed to cognitive illusions is that the illusion persists even once the subject is informed about the nature of the illusion. And the lines of the Müller-Lyer illusion continue to *look* of different lengths even once we know they are not. This may indicate that the effects elicited by optical illusions – the inappropriate deployment of conceptual-recognitional capacities – occur at what would be called a 'low' or 'subpersonal' level within the subjects' perceptual systems, i.e. via mechanisms more or less impervious to conscious thought or belief (Fish 2009: 176). The conceptual-

recognitional capacities here may indeed scarcely involve the application of anything we would call 'concepts', and may even involve innate mechanisms rather than learned facts or associations (ibid.).²⁴

There are more clear-cut cases of optical illusion in which the illusion seems less easily attributable to cognitive factors, whether at a 'low' or 'high' level. These include the checkerboard illusion (figure 2):



Figure 2 – the checkerboard illusion. The squares labelled 'A' and 'B' are the same shade (luminance) but B looks lighter because it appears to be in shadow. (Image copyright Edward Adelson: <u>http://web.mit.edu/persci/people/adelson/checkershadow_illusion.html</u>)

We might try to explain this illusion using the two-faculty account, by claiming that our experience has as its 'core' the simple presentation of the squares with their respective shades, while the illusion arises with some further operations of our visual or cognitive

²⁴ Note that, as Fish acknowledges, his fact-based account of cognitive illusion would not explain the illusory appearance in this case, since the cognitive-recognitional capacity would fail to latch onto any actual facts – the lines are, after all, not actually the corners of three-dimensional, right-angled objects (Fish 2009: 174). As I argued in the previous section, using the example of misperceiving a rope as a snake, the same consideration undermines his explanation of other cases of illusion too. It is unclear whether Fish's appeal to some 'low-level' and possibly innate mechanisms gets round this difficulty.

system, which yield a mistaken judgment about similarity of shades. If this is intended to effect the desired division of phenomenological labour then it comes up hard against our intuitive response to how the drawing looks. Here is a case where ascribing the illusory appearance to some sort of judgment or inclination to judge proves rather unconvincing. In other words, there really seems to be a different quality inhering in the image where it depicts square A than there is where it depicts square B. Certainly we are inclined to judge that A is darker than B, and indeed find it very difficult to resist that inclination, but this inclination to judge accordingly is quite compellingly grounded in a more fundamental sense in which it just *looks darker*. Perhaps what makes this illusion more compelling than the Müller-Lyer illusion is that it is easier to attend to the illusory aspect of the figure; the loci of the illusory qualities are more clearly delineated in the checkerboard case and more able to withstand close and prolonged scrutiny.

The attempt to separate 'core subjective character' from 'cognitive phenomenology' seems especially strained once we consider a chromatic relative of the checkerboard illusion (figure 3):



Figure 3 – the Rubik's cube illusion. What appear to be depicted as blue squares on the upper surface of a cube in the left-hand image are objectively the same colour (grey) as what appear to be yellow squares on the upper surface of the cube in the right-hand image. (Image copyright R. Beau Lotto: http://www.lottolab.org/illusiondemos/Demo%2012.html).

Certainly, the illusory aspects of these images are in some sense attributable to goingson internal to the subject. With both the checkerboard and Rubik's cube illusions, the illusion can be understood as the result of an adaptive response of our visual system to local and global variations in illumination. Both illusions depict real-world scenarios involving variation in illumination. The checkerboard illusion depicts the changing appearance of a surface given local variation in background illumination, while the Rubik's cube illusion depicts differences in the appearance of the same object resulting from changes in the colour of illumination (or perhaps from viewing the object through differently-coloured lenses). The illusions exploit our ability to compensate for changes in viewing conditions – that is, to disregard local variations in illumination due to shading when assessing the real lightness of objects, or to compensate for the colour of background illumination when assessing the real colour of objects (where, without committing ourselves to any particular ontological view of colour, 'real colour' can be taken to mean the colour as seen in broad daylight). Some or all of the illusory effect in these cases may reflect what is known as simultaneous contrast, a phenomenon in which a given patch of colour appears to be a different shade or hue depending on the colour of surrounding regions. It has been suggested that simultaneous contrast effects result precisely from an evolved capacity to track objects' colours through changes in environmental viewing conditions (Lotto & Purves 2000: 12836-7).²⁵

The adaptive value of these illusion-generating cognitive processes undermines one possible defence of the natural view from these sorts of optical illusion. Such a defence might claim that the internal goings-on in these cases do not make a positive contribution to phenomenology, but merely impose a negative constraint. Rather than altering or adding any phenomenal qualities they simply distort our awareness of the objects and qualities that are wholly 'out there'. One might therefore be tempted to posit a similar constraining role for the subject's physiology in optical illusions as has been suggested for cases where a subject's perceptual acuity is limited, e.g. blurry vision. While blurry vision might be presented as a problem for the natural view – where is the blurriness if not out

²⁵ The illusory aspect of simultaneous colour contrast examples has been attributed by some authors to a misjudgment: as Hamlyn notes, it might be supposed that in such cases "the perceiver makes an inference about the colour of the inset patch. That is to say that there is a tendency to infer that the inset colour must be opposed to the colour of the surrounding area, and this inference sometimes overrides the given facts of sensory experience" (Hamlyn 1983: 13). Here again we can find a version of the two faculty view: "the given facts of sensory experience" are what is supposed to be delivered by the faculty of (direct) awareness, while the 'inference' is clearly an operation of the cognitive faculty.

there? – the plausible rejoinder is that blurriness is not some extra quality of which we are aware but a limit or indeterminacy in respect of some quality of which we are aware (see Fish 2009: 54-8 for a version of this defence). A similar approach to optical illusions might be to say that some shortcomings in our visual system are such that we cannot accurately perceive certain qualities under certain conditions; for example, we cannot accurately see spatial relations such as two lines' being parallel or of equal length when they are presented in a certain confounding context. Whether or not this 'distortion defence' is plausible even for cases like the Müller-Lyer illusion, it looks highly implausible for cases like the checkerboard or Rubik's cube illusions. In part, this implausibility is intuitive: these cases just don't *look* as though they involve distortions or indeterminacies. This intuition is supported by the fact, noted above, that the illusory aspect in these cases is readily pinned down to a defined locus that can be the focus of scrutiny without dispelling the illusion. More importantly, however, the distortion defence is rendered implausible by the observation that at least some optical illusions arise from compensatory mechanisms in the visual system that have adaptive value. Indeed, the illusions are side-effects of cognitive processes that, in normal environments, enhance the acuity of vision in respect of the subject's ability to identify and discriminate objective surface qualities of objects despite potentially confounding variations in illumination, etc.

What we have then are cases of optical illusion in which the illusory aspect seems to involve genuine aspects of phenomenology that are not attributable to the worldly objects or scenes perceived, or to the subject's (spatial) perspective on them. Further evidence for the positive role of internal processes in determining phenomenology comes from a different sort of illusion, in which one is aware of phenomenal qualities of a *kind* that is not exemplified in the object perceived. The illusions concerned involve motion-induced colour, in which an object lacking hue (being black and white) appears, under a certain motion, to exhibit hues. A classic example is Benham's disk, in which a patterned black and white disk is spun rapidly while being viewed. The subject as a result sees flickering colours in the rotating disk. This phenomenon clearly meets the criteria for being an optical illusion, being both predictable and repeatable, and not reliant on the subject's background beliefs. What's more, the colours are clearly perceived as qualities of the object seen – as located on the surface of the disk – and so would not be aptly described as hallucinatory phenomena somehow induced by the disk but 'floating free' of it, as it were. Thus, with motion-induced colour we have an illusory phenomenon in which phenomenal qualities appear to be positively *added* to the experience and not merely altered in some way (as we find in the checkerboard and Müller-Lyer illusions for example). This therefore poses the starkest challenge yet to the natural view (with hallucination waiting in the wings, as the subject of the next chapter).

We might nonetheless find a role for the simple metaphysical picture in explaining these more intractable cases of illusion. After all, the simple metaphysical picture is explicit that experience is constituted by the whole causal perceptual process in which the object affects the subject, a process that is partly constituted by events internal to the subject. Why is the problem, then, with allowing that these internal events are partial determinants of phenomenology? The problem is that this risks neglecting the simple metaphysical picture's role in underpinning the natural view. We could just assert that the experience is identical to the process and that the phenomenology is itself the whole process or an aspect of the process, in which case it is not unreasonable that altering one part of the process, even if physically internal to the subject, might change the overall phenomenology. This, recall, is precisely the account offered by Manzotti in his 'process-oriented view'.²⁶ But recall also that Manzotti's account was rejected precisely because it conflicts with the natural view since, even in cases of veridical perception, it fails to locate perceived qualities where they seem to be, namely out in the subject's environment.

There are further features of ostensibly *veridical* perception that put similar pressure on the natural view. These include variability in colour perception among different normal subjects, and even the nature of colour itself. I will consider these seeming problem cases in chapter 6 below, before suggesting how the simple metaphysical picture might offer a solution after all.

²⁶ See Manzotti (2006a: 29-30) and (2008: 186-7) for his explanation of illusion. His (2006a) and (2008) address optical and physical illusions respectively.

3.6 Conclusion - the argument(s) from illusion

From our survey of illusion we are left with the following conclusions:

- Physical illusions present no difficulties for the natural view, being wholly explicable by reference to the nature and qualities of the objects perceived, their spatial relation to the subject and the background conditions of perception (illumination etc.).
- 2. Cognitive illusions may in some instances be explicable via the two faculty view, leaving the phenomenology accounted for by the objects and qualities perceived, and the illusion explained in terms of a tendency towards misjudgment about the nature of those objects or qualities. There are other cases, however, in which the cognitive faculty appears to play a genuine role in influencing experiences' phenomenology (Brewer's 'core subjective character'); these include cases of cognitive penetration of perceptual experience, which I will discuss in chapter 5 below.
- 3. Optical illusions may also in some instances be explicable via the two faculty view. Again, however, there are cases in which the illusion seems to reside not in a judgment but in the very qualities of which we are aware. These included cases in which we seem to be aware of certain qualities (hues) that seem not to be objectively present in the appropriate part of our environment (the Rubik's cube illusion, motion-induced colour).

At the end of this chapter we are therefore left with a number of cases that are not readily assimilated to naïve realism and its explanatory virtue of accounting for the phenomenology of experience by simple appeal to the antecedent nature of the objects perceived. I suggested that Fish's fact-based account of cognitive illusions fails to account for our intuitions in some cases, leaving us with a residual worry that our conceptual-recognitional capacities might play a positive role in influencing the phenomenology of experience even when they are deployed erroneously. As it is, I will describe in chapter 5 cases of perceptual learning and cognitive penetration of perceptual experience which provide strong evidence for the role of memory and belief in positively influencing the phenomenology of perceptual experience. I will show, however, that the natural view, and

naïve realism more specifically, can accommodate such cases using a mechanism made possible by the simple metaphysical picture of experience.

Some cases of optical illusion, such as the Rubik's cube illusion, will prove more intractable. There appears to be no solution to these cases that can be accommodated to naïve realism. In chapter 6 I will describe other features of perceptual experience that pose a similar challenge. I will argue that, again with the help of the simple metaphysical picture, we can account for these consistently with the natural view, but only at the cost of rejecting the realism claim that is essential to naïve realism. In other words, such features of experience demonstrate that the worldly objects and qualities of which we are aware are in some respects not constitutively independent of the subject or her experience. This will seem a high price to pay to preserve the natural view, but I will in chapter 7 argue that it can be accommodated within a broadly physicalist metaphysics and so avoid the charge of idealism.

CHAPTER 4 - PHENOMENOLOGICAL PROBLEMS (II): HALLUCINATION

4.1 Introduction: hallucination and the natural view

We have seen that at least some illusions can be accommodated to the natural view and its insistence that phenomenal character is to be found 'out there', borne by worldly objects. That is as well, since the natural view is a claim about perception and illusions are perceptual even if they are misleading. A different sort of sensory (or quasi-sensory) experience has widely been seen to present even greater challenges to naïve realism and, by extension, the natural view. That is *hallucination*. Since hallucination is, by definition, not a form of perception, it is not directly a challenge to the natural view. Instead, as I explained in chapter 3, the challenge comes indirectly via the thought that whatever accounts for the phenomenology of hallucinatory experience ought equally to account for the phenomenology of perceptual experience. I will examine this 'argument from hallucination' in more detail in the following section.

As we shall see, naïve realists have typically responded to the argument from hallucination by denying that the phenomenology of perceptual experience and that of hallucination need to be explained in the same way. I will take a quite different approach, arguing that the simple metaphysical picture of perception allows us to accept that the phenomenology of perception and hallucination *is* best given a single kind of explanation, but that this need not undermine the natural view of perception. Specifically, we will be able to offer this unified account by attributing the hallucinatory phenomenology to awareness of worldly objects. Accepting this admittedly perverse-sounding claim will require a reworking of our definition of hallucination in order to allow that hallucinations too can have worldly objects in a certain fashion.

Before reaching such a counter-intuitive conclusion, I will consider some alternative strategies proposed by naïve realists to defend their theory against hallucination. I will argue that these alternatives either fail to account positively for the phenomenology of hallucination (some even denying that hallucinatory experience *has* phenomenology) or fail to account for what is distinctive about the phenomenology of perceptual experience that makes it uniquely world-presenting.

The account offered here, and based on application of the simple metaphysical picture, acknowledges that hallucinatory experiences involve genuine awareness of the same sorts of qualities that we can become aware of through perception. Furthermore, it explains that the phenomenology of hallucination is had by virtue of the subject's standing in an awareness or acquaintance relation to worldly objects previously perceived. Counter-intuitive as this account might seem, I hope to show that it can be made plausible by a four-stage argument which, firstly, treats hallucinations as dependent on the subject's memory of things previously perceived; second, takes such perceptual memories to involve a kind of 'delayed' awareness; third, appeals to intermediate cases to persuade that a sharp distinction between perception and memory is not defensible; and, fourth, takes this as justification for extending the simple metaphysical picture of perception to cases of perceptual memory and, thereby, hallucination.. Crucially, the result is an account with much greater explanatory power than those existing naïve realist theories which either deny that hallucinations possess genuine phenomenology or attribute it to awareness of entities other than worldly objects.

4.2 The argument from hallucination

Recall the argument from hallucination:

The argument from hallucination:

- 1. For every perceptual experience there could be a hallucinatory experience that seems the same to the subject.
- 2. The best explanation for the fact that two experiences seem the same is that they have the same metaphysical nature.
- 3. All hallucinatory experiences involve awareness of phenomenal qualities that are not instantiations in worldly objects; those qualities are not 'out there', inhering in worldly objects.

Therefore

4. All veridical perceptual experiences involve awareness of phenomenal qualities that are not instantiations in worldly objects; those qualities are not 'out there', inhering in worldly objects.

As noted earlier, this key step in this argument is premise 2, an inference to the best explanation, such that two experiences that seem the same ought to share a metaphysical nature. The argument thus explains a phenomenological equivalence between possible perceptual and hallucinatory experiences – their subjective indistinguishability – as being the upshot of a metaphysical equivalence; they seem the same because they *are* the same. If we further accept that hallucinatory experience is not a relation of awareness to worldly objects, a claim (premise 3) which is apt to seem uncontroversial, it follows that perceptual experience is also not a relation to worldly objects. A key feature of the argument from hallucination is thus its satisfaction of an underlying urge towards explanatory unity: whatever explains the phenomenology of hallucinatory experience should also explain the phenomenology of perceptual experience, and if it's not the constitutive role of worldly objects in one case it should not be in the other case either.

As we shall see, standard naïve realist counter-arguments attack precisely the move from phenomenological to metaphysical equivalence and so leave open the possibility that perceptual and hallucinatory experiences, even though they might *seem* the same, are in fact quite different *kinds* of experiences in some metaphysical sense. This in turn allows the naïve realist to hold onto the claim that perceptual experience is genuinely relational and involves direct awareness of worldly objects and their qualities, even if the same is not true of hallucinatory experience. Rejecting the move from phenomenological to metaphysical equivalence does, however, leave those forms of naïve realism without the benefit of the explanatory unity that is claimed by the proponents of the argument from hallucination. In other words, the naïve realist is left having to explain how two experiences – one involving awareness of worldly objects and the other not – might come to seem the same despite their radically different metaphysical natures.

Thus, while one could reject the argument from hallucination, and so defend the natural view, by denying any or all of premises 1, 2 or 3, the standard naïve realist counterarguments have focused on the denial of premise 2. All agree that hallucinations lack worldly objects (premise 3) while denying that matching hallucinatory and perceptual experiences share a metaphysical nature. Premise 3 is apt to seem incontrovertible – it is often taken to be definitional of hallucinations that they lack worldly objects – so it is no surprise that naïve realists have almost invariably attacked premise 2 instead.

By sharp contrast, I will take the opposite course to defend the natural view, by offering an alternative account, based on the simple metaphysical picture, which accepts premise 2 while accepting the truth of premise 3. The key claim is, counter-intuitively, that the phenomenology of hallucinatory experiences is wholly attributable to awareness of worldly objects and their properties, a claim which can be rescued from absurdity by careful application of the simple metaphysical picture. It follows from this view that premise 2 is true by virtue of a restricted range of possible cases in which awareness of the same scene yields the phenomenology of both perceptual and hallucinatory experiences.²⁷

Before considering this and the other counter-arguments to the argument from hallucination, let us look in more detail at the argument itself. Premise 1 makes the plausible claim that, for any given perceptual experience, there could be a hallucinatory experience which is exactly matching subjectively, so that the subject could not tell by reflection alone that the experience was hallucinatory rather than perceptual.²⁸ The move from phenomenological sameness to metaphysical sameness is made explicit in the move from premise 1 to premise 2, and shown to be an inference to the best explanation. What it is for two experiences to be matching is their possessing the same phenomenology, and this matching phenomenology is attributable to a 'match' in the metaphysical structure of experience across the perceptual and hallucinatory cases, namely that in each case the subject stands in some awareness relation to worldly objects and their qualities.²⁹

²⁷ Given the account I am about to offer, the obvious kind of case would be one in which a subject hallucinates the same scene that she perceived at some earlier time (where its being a hallucination of the same scene I take to mean that it involves awareness of the same scene, with its constituent objects and qualities).

²⁸ As noted earlier, the requirement that subjects be unable to distinguish experiences 'by reflection alone' is introduced by Martin to avoid the possibility that subjects could tell that they were perceiving and not hallucination (or *vice versa*) because they were told as much or in virtue of their having some pertinent background beliefs (Martin 2006: 364-5). For example, a subject might deduce that she was merely hallucinating her best friend before her if she knew that the friend was currently on vacation on the other side of the world.

²⁹ This is often expressed as the claim that matching hallucinatory and perceptual experiences share the same phenomenal character, although this typically operates within an understanding of 'phenomenal character' as something possessed by experiences potentially independently of what the experiences are *of*, i.e. a 'common factor' shared by indistinguishable experiences. As noted earlier, naïve realists often offer a metaphysical construal of 'phenomenal character' which takes it to be constituted by the worldly objects perceived and their qualities. On my use of 'phenomenal character', meanwhile, it means that relational property of scenes which is how they appear to a subject, which allows that different scenes can

It now looks as though what makes an experience *perceptual* cannot be its phenomenology – for this is shared with the matching hallucination – but some other factor that connects the phenomenology with the worldly object perceived. Generally, this extra factor is understood to be causal. Perceptual experience is thus taken to be analysable into two elements: first, the experience itself with its phenomenology; and, second, a causal link that makes the experience an effect of the object(s) perceived. This amounts to an endorsement of the causal theory of perception and the existence of a common factor amongst matching perceptual and hallucinatory experiences (see chapter 2).

Central to most defences of naïve realism is a rejection of the inference from a phenomenological sameness among experiences to their metaphysical sameness. Since this inference is merely abductive, rejecting is not in itself too difficult – the metaphysical and phenomenological construals of 'matching' are neither equivalent nor mutually entailing. There is an entailment in one direction – experiences that are matching in the metaphysical sense will be matching in the phenomenological sense – but it does not follow from the fact that two experiences are matching in the phenomenological sense that they have the same metaphysical nature.

Even lacking this latter entailment, the explanation of phenomenological sameness via metaphysical sameness might seem compelling as an inference to the best explanation. Note that the inference does not rely on assuming that the phenomenological sameness is due to a sameness in the arrangement of objects and qualities of which the subject is aware. As Austin points out, two objects might be quite different and yet look the same – a lemon and a yellow, lemon-shaped bar of soap, for example (Austin 1962: 50). The posited metaphysical sameness instead claims that the perceptual and hallucinatory experiences have, more broadly, the same metaphysical structure, e.g. involve a relation of awareness between subjects and the same kinds of objects or properties, whether these are considered to be 'out there' or 'in here'.

If we accept that the phenomenology of hallucination is not attributable to awareness of worldly objects and qualities (as per premise 3), then the natural view forces us to

bear the same phenomenal character. I will therefore allow that experiences of different scenes can present or exhibit (not possess) the same phenomenal character.

conclude that perceptions and hallucinations can be subjectively matching despite some profound difference in the nature of the objects of awareness (if any) and so a corresponding difference in the overall metaphysical structure of the respective experiences.

A naïve realist who rejects the move from epistemic sameness to metaphysical sameness therefore takes on the burden of explaining how two very different kinds of phenomenon could nonetheless seem the same. The first step is to be clear about the metaphysical difference that is required. According to the natural view, the phenomenology of perceptual experience consists of worldly *objects* with their perceptible qualities, and not merely the qualities alone. Accordingly, naïve realists have typically responded to the argument from hallucination by denying that perceptual and hallucinatory experiences could share the same phenomenology. The denial is made precisely on the grounds that the phenomenology in perception is attributable to worldly object, while in hallucinatory experience there is no corresponding worldly object. It is sometimes said, on this basis, that the perceptual experience is 'object-involving', unlike its hallucinatory counterpart (Snowdon 1990: 130).

As noted in chapter 2, proponents of this defence of naïve realism are termed 'disjunctivists' because they advocate a disjunctive analysis of the notion of sensory experience. Thus, its appearing to me that there is a red tomato might be made true by one of two possible states of affairs: (a) there being something that looks to me to be a red tomato; or (b) its merely seeming to me as if there is something that looks to me to be a red tomato (cf. Snowdon 1981: 185). Note that in the first, perceptual, disjunct, it need not be the case that there is in fact a red tomato before me for this to count as a case of perception – what matters is that there is an object that I am seeing and which I take to be a red tomato. This allows that illusory experiences are also cases of perceiving, albeit misleading ones.³⁰

³⁰ This form of disjunctive analysis thus places veridical and illusory perceptual experiences together, a version of disjunctivism that Byrne and Logue call 'VI v. H disjunctivism' (2008: 69). As they note, an alternative form of disjunctivism has been proposed which places veridical perception in one disjunct and illusion and hallucination in the other ('V v IH disjunctivism'). A clear example of this is the disjunctivism proposed by John McDowell, but which has different motivations from those claimed in the current thesis, namely a concern to defend the special epistemic value of veridical as opposed to subjectively indistinguishable but non-veridical experiences (i.e. illusions and hallucinations) (see e.g. McDowell 1994).

The disjunctive analysis is intended as an alternative to what has historically been the standard view, namely that a given statement of how things seem sensorily to a subject will be made true in all cases by the occurrence of a certain experience, where this is conceived as, in some sense, an inner state or event within the subject. This standard view is one in which experiences are individuated by their phenomenology, and phenomenology is in turn identified with what it's like to undergo that experience. What, on the standard view, distinguishes perceptual from non-perceptual experiences is not the experience itself but the aetiology of the experience – its being appropriately caused by a worldly object, or not. Since it posits a factor common to both perceptual and non-perceptual experiences – namely the experience itself – this is sometimes called the 'common factor view' of experience (Fish 2009: 3-5). And since it analyses perception into two independent conditions – the occurrence of an experience plus the obtaining of an appropriate worldly cause – it is also described as 'conjunctivism' (Johnston 2004: 114).

We can therefore see disjunctivism as the rejection of one line of thought that leads from a phenomenological equivalence between experiences to a metaphysical equivalence – the idea that two sensory experiences that seem alike must thereby *be* alike. By incorporating the worldly object within the experience itself (in the perceptual but not the hallucinatory case) disjunctivists undermine this distinction and so undo the entailment from phenomenological to metaphysical sameness.

Having abjured the explanatory unity offered by the argument from hallucination, the disjunctivist is left seeking an explanation of how perceptual and hallucinatory experiences might seem the same despite their profound metaphysical differences. There might be broadly two routes for the disjunctivist to follow here, which we can bring out with the following observation: it seems to the perceiving subject that she is aware of worldly objects and their qualities, and it seems that way because she *is* so aware; it likewise seems to the hallucinating subject that she is aware of worldly objects and their qualities, but in this case there are grounds for thinking she is *not* so aware. If we therefore say that, in the hallucinatory case, the subject merely *seems* to be aware of worldly objects and qualities, we leave some ambiguity in what this *seeming* amounts to. Does it mean that the subject is not really aware of objects and/or qualities? This question points to two broad

approaches to a disjunctivist account of hallucination, both of which are represented in the literature:

- Hallucinatory experience involves awareness of qualities, but these qualities are not instantiated in worldly objects of which the subject is aware.³¹
- 2. Hallucinatory experience does not in fact involve awareness of qualities at all, but merely *seems* to do so.

As stated, option 1 leaves open the question of where the hallucinated qualities are instantiated, and so potentially paves the way for different versions of disjunctivist naïve realism. In fact, as we shall see, it leaves room for the claim that hallucinated qualities are not instantiated in *any* objects, worldly or otherwise. Option 1 can therefore be subdivided as follows:

- 1a Hallucinatory experience involves awareness of qualities but these qualities are instantiated in something other than worldly objects
- 1b Hallucinatory experience involves awareness of qualities, but these qualities are not instantiated in any objects, worldly or otherwise, of which the subject is aware.

I will discuss each of these options in turn, assessing their coherence and their capacity to explain how hallucinations and perceptions can seem the same while having a different metaphysical structure. I will conclude that none offers a satisfactory explanation in this regard, leaving us with either a glaringly disunified explanation of the subjective similarities among different experiences (option 1a), an appeal to metaphysical oddities (options 1b and, in some forms, 1a), or an unsatisfying denial that hallucinations even involve awareness of qualities in the first place (option 2).

³¹ When referring to hallucinatory experiences I assume for the sake of exposition that they are 'total' hallucinations i.e. the experiences are entirely hallucinated and involve no perceptual element.

With these shortcomings in mind, I will go on to propose a further option, one that takes the unusual step of denying premise 3 of the argument from hallucination rather than premise 2. In other words, it claims that hallucination does in fact involve the awareness of qualities instantiated in worldly objects, just as perception does:

 Hallucinatory experience involves awareness of qualities that are instantiated by worldly objects, just as they are in genuinely perceptual experience.

I will argue that option 3, although seemingly absurd, can be made intelligible by appeal to the simple metaphysical picture and in light of a plausible account of the aetiology of hallucination. Allowing that hallucination is object-involving in the relevant sense – being capable of description as a process relating the subject and worldly objects – has the great advantage over the other options of yielding a unified explanation of phenomenology across both perceptual and hallucinatory cases. Otherwise, the choice between options 1a and 1b, on the one hand, and option 2 on the other, presents something of a dilemma: either we explain hallucination as awareness of qualities not instantiated in worldly objects, or we deny the 'felt reality' of hallucination in the first place. The first horn of the dilemma threatens to undermine the natural view and its privileging experience as a way of directly encountering the world; while the second calls into question our intuitive grasp of the subjective, quasi-sensory character of hallucination.

Option 3 rescues us from this dilemma by claiming that in hallucinatory experience we are aware of actual qualities, and that these qualities are instantiated in worldly objects. Implausible as this seems, the simple metaphysical picture offers a way to make it intelligible. The account has great explanatory virtues. By holding that hallucinations, like perceptual experiences, are object-involving, it offers to explain the phenomenology of hallucination in terms of things we already understand (worldly objects and processes). What's more, it offers a single, unified account to cover both perceptual and hallucinatory experiences.

Before laying out option 3 in more detail, let us consider each of options 1a, 1b and 2 in turn, noting their advantages and, more importantly, disadvantages.

4.3 Option 1a - awareness of non-worldly qualities

The first option might, for example, involve the claim that the phenomenal qualities of hallucinatory experience are qualities of the experience itself, i.e. something like the qualia described earlier. Or it might involve the claim that the phenomenal qualities are instantiated in mind-dependent objects – sense-data, ideas or impressions.

It is not easy to find clear examples of such a 'mixed' theory of phenomenology – one in which the phenomenology of perceptual experience is determined or constituted by worldly objects and their qualities, while the phenomenology of hallucination is determined or constituted by non-worldly objects or qualities. One example might be recent defences of the 'theory of appearing' (Alston 1999; Langsam 1997). According to the theory of appearing, perceptual experience is to be understood as a relation holding between subject, object(s) and some qualities. The relation is one of appearing (looking, sounding, etc.). So, for example, a subject *S*'s visual experience of a ripe tomato could be described as 'The tomato's looking red to *S*'. Insofar as this characterisation makes the object perceived a constituent of the experience, the theory of appearing stands as a form of relationalism as described in chapter 1. And, depending on how one uses this relational claim to explain phenomenology, the theory of appearing could be construed consistently with the natural view.³²

As a form of relationalism, the theory of appearing faces the same problem from hallucination as any theory endorsing the natural view. Since it explains perception, and thereby phenomenology, as a relation to worldly objects, how will it explain phenomenology in cases where a worldly object is not perceived? Both Langsam and Alston endorse a version of disjunctivism: they point out that, while perceptual and hallucinatory experiences might appear the same, this does not entail that they *are* the same, metaphysically speaking (Langsam 1997: 39; Alston 1999: 190). Both presume that hallucinations can involve awareness of qualities, but deny that these need to be the same sorts of qualities, or inhere in the same sorts of objects, as those qualities we are aware of

³² Although Langsam describes the theory of appearing as a form of naïve realism, his account falls short of endorsing the natural view, identifying phenomenal qualities with *extrinsic* properties of perceived objects; specifically, their properties of *appearing F* (red, sweet, etc.) to the subject (Langsam 1997: 53-56).

when we perceive. No fully worked-out alternative is offered; both consider noncommittedly the notion that hallucinatory experiences might be relations to regions of space – the regions that the hallucinated objects seem to occupy (Langsam 1997: 47; Alston 1999: 191), although the obvious shortcoming here is that it tells us only where the (metaphysically distinct) phenomenal qualities of the hallucination seem to reside, and not what they are or how they come to feature in experience in this way. Alston favours instead the idea that hallucinations involve awareness of mental images, although he does not claim to offer a robust exposition or defence of the suggestion (Alston 1999: 191-2).

In any case, the worry with any such 'mixed' account is just the one we noted at the outset: it offers a glaringly disunified explanation of phenomenology. Sometimes this is attributable to awareness of worldly objects and qualities, and sometimes to awareness of qualia, mental images or mind-dependent objects, etc. We can intuitively understand how qualities such as colour can belong to worldly objects (our intuitive commitment to this is, after all, part of the motivating insight behind the natural view), but it is difficult to grasp how they can also belong to such curious entities as 'experiences' or 'sense data'.

Such views might also seem to undermine the insistence that genuinely perceptual experience involves direct awareness of worldly objects. What is distinctive about such awareness if a subjective similar hallucinatory experience can result from awareness of something else – sense data, mental images, etc.? There might remain epistemological motivations for insisting on direct awareness of worldly objects – for example, to explain our ability to think about particular worldly objects (see e.g. Campbell 2002a: 114 ff.) – but the phenomenological motivation lying behind the natural view would seem to be much weakened. This is, of course, just to acknowledge the urge to unifying explanation which drives the argument from hallucination. As we shall see, the same worry afflicts the next option.

4.4 Option 1b - uninstantiated universals

This option is proposed by Mark Johnston, an avowed 'radical direct realist' who suggests that hallucinatory experience involves the awareness of a 'sensible profile' – a structured arrangement of qualities. The sensible profile is just the same as that which a subject might be aware of in genuinely perceiving a scene – it is after all just the array of

qualities that determines what it is like to perceive that scene; the yellowness of the autumn foliage against the blue of the sky perhaps. Crucially, though, a corresponding hallucination of autumnal trees would, on Johnston's view, involve awareness of the same kind of sensible profile but one that is not instantiated in worldly particulars or indeed in any objects, worldly or otherwise (Johnston 2004: 134 ff.).³³

This account is unlike Option 1a in that it is not a 'mixed' view – it does not postulate different kinds of objects to be the bearers of phenomenal qualities in the perceptual and hallucinatory cases. Indeed, it posits a common factor among these cases – the sensible profile shared by subjectively indistinguishable perceptual and hallucinatory experiences, awareness of which determines what it's like to have those experiences. To that extent, it departs from the disjunctivist rejection of a common factor shared by perceptual and hallucinatory experiences and which is supposed to determine the experiences' phenomenology (Johnston 2004: 170-1).³⁴

Option 1b thus has an explanatory virtue inasmuch as it offers a unified account of phenomenology in both kinds of experience. However, its explanatory power proves a somewhat mixed blessing. By giving a positive explanation of phenomenology in the hallucinatory case, in terms of uninstantiated sensory profiles, the account offers something that might do the explanatory work for the perceptual case too. Thus, if sensory profiles constitute the phenomenology of the hallucinatory cases, why not think that they constitute the phenomenology of the perceptual cases too? In other words, if veridical perception and hallucination can share the same sensible profile, and this sensible profile constitutes what it is like for the subject to have the experience in question (constitutes its phenomenology, in other words), then there doesn't seem to be any role for the particular *qua* particular (as opposed to the instantiation of a sensible profile) in determining the phenomenology of the experience in the veridical case. There is, we might say, no way in

³³ It is important that perception, for Johnston, involves the awareness of *instantiations* of sensible profiles and not merely sensible profiles *that are instantiated*. As Dunn points out, there is nothing to rule out a particular sensible profile's being instantiated even when it is the primary object of a hallucination, since I might hallucinate a sensible profile which is not instantiated by anything in front of me (as of, say a spiny anteater) but which *is*, coincidentally, instantiated on the other side of the world. If veridical perception were construed as the awareness of sensible profiles that are instantiated (somewhere), this would have the absurd consequence that I am veridically perceiving a spiny anteater (Dunn 2008: 379).

³⁴ Johnston is explicit that his theory is non-disjunctivist insofar as it admits a common mental factor amongst perceptual and hallucinatory experiences, namely the (awareness of the) sensible profile that is instantiated in the former case but not in the latter. However, it remains disjunctivist to the extent that it allows for an experiential report to be made true by two different kinds of state: one that is awareness of an instantiated sensible profile, and another that is awareness of an uninstantiated sensible profile.

which the worldly objects can 'show up' in phenomenology over and above what is provided by the sensory profile. As Martin puts it, the feature that determines phenomenology in the hallucinatory case 'screens off' the 'naïve realist aspects' from making any distinctive contribution to the phenomenology in the perceptual case (Martin 2004: 62).

To be fair, Johnston is at pains to emphasise that the subject is not, in perceiving, aware of the object *in virtue* of being aware of the sensible profile, for that would make object perception indirect for his purposes:

[I]t is not that we see particulars *by* being visually aware of the sensible profiles that they have or instantiate. Instead, we see them *as* having certain sensible profiles, which are after all just certain complex sensible profiles. We may have to be visually aware of those sensible profiles in order to see particulars as having them, but this implies no indirection.

Johnston 2004: 155

Not only does this avoid implying any 'indirection' (the sensible profile as an intermediary object of awareness), but it also signals that, first and foremost, what is perceived is the *object* which bears the relevant sensible properties.

If there is a residual worry with the sensible profile account, however, it is that its postulation of awareness of uninstantiated sensory profiles undermines any phenomenological motivation for thinking that in genuine perceptual experience our awareness of sensible profiles is thereby a direct awareness of the *objects* that instantiate them. Once we have admitted awareness of uninstantiated sensory profiles it seems that direct awareness of objects has at best a diminished role to play in the explanation of phenomenology. If a given sensible profile could be common to both a perceptual and a hallucinatory experience, what grounds do we have for imputing a direct awareness of objects in the former case?³⁵

³⁵ Johnston's account is therefore vulnerable to the 'screening-off' problem raised by Martin, namely the worry that any common factor shared by hallucination and perception, and which is responsible for the phenomenology of hallucination, will inevitably 'screen off' the actual worldly objects from having any phenomenological impact in the case of perception (Martin 2004: 62). I will discuss the screening-off worry further in section 4.5 below.

A similar objection faces any theory that offers a positive explanation of phenomenology as awareness of anything other than worldly objects. The natural view is motivated by the intuition that perceptual experience is a direct apprehension of worldly objects so that its phenomenology is simply composed of those objects. If phenomenology might just as well be composed of non-worldly objects or uninstantiated qualities then worldly objects lose their special or distinctive role in constituting phenomenology, and the motivating intuition behind the natural view is undermined.

Behind this worry lies what we might see as the core weakness of Johnston's theory: its offering a positive explanation of the manifest character of hallucination in terms of something – the (uninstantiated) sensible profile – that itself needs considerable explanation. The appeal of naïve realism, and the natural view in particular, is its explanation of phenomenology in terms of familiar worldly objects, but this appealing feature is undermined by the introduction of such mysterious entities as uninstantiated sensible profiles.³⁶ Any explanatory virtue the view might claim (the appeal to a common explanation of phenomenology) thus needs to be set against a considerable metaphysical vice. Compared with option 1a, the account has the virtue of not appealing to such obscure entities as sense-data or qualia, but it gains this by introducing the equally puzzling notion of 'free-floating', uninstantiated qualities, presumably to be given a special ontological status as universals.

Sensible profiles might seem candidates for what others have called sense-data, although Hilbert suggests that Johnston's view is better understood as a version of representationalism (Hilbert 2004: 190). As such, the suggestion that experience involves the awareness of sensible profiles might be taken as a gloss on the claim that experience *represents* the instantiation of sensible properties. According to certain 'particularist' forms of representationalism, perceptual experience has representational content that includes as constituents the mind-independent objects perceived along with their properties. The content is 'singular' in taking the perceived particular as a constituent, but the content also includes the particular's properties. One option available to the particularist representationalist when challenged to account for hallucinatory experience is to claim that such experience also has representational content but the content is *existential* rather than singular, i.e. it represents that *there is something that is F* rather than that that to is F. Another way to put it is to say that the representational content of experience has one or more 'slots' for the particular(s) perceived, that these slots are filled in the case of perceptual experience but 'gappy' when one hallucinates (Tye 2009: 80-83). The upshot of this claim – which Tye calls the 'Singular (When Filled) Thesis' - is that the phenomenology of perceptual and hallucinatory experiences can be composed of the same qualities, namely those that are represented in the existential content. What distinguishes them is that the content of the perceptual experience is singular, while that of the hallucinatory experience is gappy (ibid.). Or, in Johnston's terms, the perceptual experience involves awareness of an instantiated sensible profile (i.e. the object that bears those sensible properties) which the hallucinatory experience involves awareness of an uninstantiated sensible profile. I don't propose to deal with this version of representationalism in detail, except to note that it suffers the same weaknesses that have been identified for Johnston's account (inevitably so, if they are structurally the same). It is also far from clear whether any perceptual theory in which the notion of *representation* is fundamental could underwrite the natural view as I have presented it.

The mystery is not dispelled by Johnston's elaboration of his metaphysical account, in which he posits what I have called the simple metaphysical picture as an account of perceptual experience:

[T]he relation between seeing an object and the long physical process involving first the light coming from the object and then the operation of the visual system is *not* the relation between a first mental effect and a prior physical process that causes it. Seeing the object is not the next event *after the visual system operates*. Seeing the object is an event materially constituted by *the long physical process* connecting the object seen to the final state of the visual system. Seeing the object is an event that is (as it actually turns out) constituted by a physical process that goes all the way out to the object seen. [...] The failure to understand the relation between the underlying causal process and seeing as material constitution, rather than process causation, is one of the deepest sources of the Conjunctive Analysis.³⁷

Johnston 2004: 138-9

He offers this statement of the simple metaphysical picture for the same reason that I use it to underpin the natural view, namely to make sense of the intuition that worldly objects 'get into' experience as constituents of phenomenology. Otherwise, as Johnston acknowledges, his claim that perceiving involves direct awareness of *more* than one is aware of in hallucinating appears at odds with "a common picture of psychophysical causation", namely the causal theory of perception discussed in chapter 2. It is to avoid the consequences of the causal theory of perception and so secure the directness of perception that Johnston offers his version of the simple metaphysical picture.

Note that Johnston applies the simple metaphysical picture only to genuinely perceptual experience: whereas the *perceptual* experience is materially constituted by the whole causal process linking subject and object, he states that the *hallucinatory* experience is materially constituted by the state of the subject's visual system alone (ibid.: 139). The reasoning behind this divided metaphysical account is obvious if we assume that there is an object of awareness only in the perceptual case. Nonetheless, this leaves us with the puzzle of how a mere state of the visual system could constitute awareness of a sensible profile. With the simple metaphysical picture we have a straightforward explanation of how objects and their qualities feature in the phenomenology of perceptual experience –

³⁷ Johnston's wording here is potentially misleading: on one reading it is uncontroversial that *seeing* is constituted by the long physical process appropriately linking subject and object, but what he means by 'seeing' here is evidently *visual experience*.

just by *being* the constituents of that phenomenology. In other words, what we call phenomenology just is the assemblage of worldly objects as they are arrayed before the subject, and which are constituents of the perceptual process that is identified with the experience itself. What we lack, in the sensible profile account, is a comparable explanation of how uninstantiated universals come to constitute or feature in the phenomenology of hallucinatory experiences. On Johnston's account, the experience is constituted by the state of the subject's perceptual (visual) system alone. There is no place for uninstantiated universals in this metaphysical picture, so positing their constitutive role in phenomenology here seems *ad hoc*. Not only are we given no explanation of why a given state of the subject's visual system would give rise to awareness of a sensible profile, we are given no explanation of why it should be awareness of this sensible profile and not any other. Perhaps the claim might be that a given neural state gives rise, in hallucination, to awareness of whatever sensible profile would normally induce that kind of neural state through perception. But in that case we would be left with the claim that neural states are sufficient from having an experience with the appropriate phenomenology, a claim that leaves genuine object awareness without a role to play in that phenomenology.

4.5 Option 2 - the no phenomenology view

We have reasons, then, to reject the first broad naïve realist approach to explaining hallucination; that is, the claim that hallucination involves awareness of phenomenal qualities that are either instantiated in non-worldly objects (sense-data perhaps) or uninstantiated or free-floating in some sense (Johnston's sensible profiles). The second option is to claim that when we hallucinate there are in fact no phenomenal qualities of which we are aware – it merely *seems* to us that we are aware of phenomenal qualities. On this account, to the extent that we can talk of hallucinatory experiences as having phenomenology this is not what we might call *robust* phenomenology (cf. Fish 2009: 83), insofar as it does not involve awareness of actual qualities, whether free-floating or instantiated in one or other kind of object. Rather, the phenomenology of hallucination is taken to be merely apparent.

The obvious advantage of this approach, which we might call the *no phenomenology view*, is that it entirely removes any pressure to give a common account of phenomenology between the perceptual and hallucinatory cases; since there *is* no phenomenology in

hallucination. In fact, it leaves the notion of hallucinatory 'phenomenology' wholly parasitic on that of the genuine phenomenology of perceptual experience: what it's like to hallucinate is not explicable in terms of any actual qualities manifest in the hallucinatory experience, but solely in terms of *what it would be like* to have a certain perceptual experience.

There are broadly two versions of the no phenomenology view. The first is developed by Martin (2004, 2006), and offers a purely epistemic account of hallucination as the subject's being in a state that she cannot know by reflection is not a veridical perception. The second is described by Fish, and explains hallucination as a mental state that has the same cognitive effects – produces the same beliefs and judgments – as a possible veridical experience (2009: 93 ff.). The two theories are motivated by the same desire to defend naïve realism from the argument from hallucination. However, although they are closely related and superficially very similar, I will note some significant differences. Ultimately, however, both theories embrace the same counter-intuitive position and so suffer the same weakness, namely that they deny our compelling intuition that hallucinations have a substantial quasi-sensory phenomenology. In other words, they are vulnerable to the challenge that they fail to account for the 'felt reality' of hallucinatory experiences. I will argue that this challenge is a powerful and compelling one, despite claims by both Martin and Fish that their respective accounts can indeed account for the felt reality of hallucination.

Let us consider first Martin's 'negative epistemic account' of hallucination. Whereas the theories considered so far have given a positive metaphysical explanation of hallucinatory experience – explaining it in terms of the nature of the objects of awareness in hallucination – Martin's account eschews all such metaphysical claims and instead asserts that hallucination is fully explicable in terms of certain limitations on our first-person, reflective knowledge of our own mental states. On this view, to hallucinate is to be in a state that one could not know by reflection alone is not a state of perceiving (Martin 2004: 76). Thus, it is a 'negative epistemic account' because it is cast entirely in terms of what a subject *cannot know*. Since Martin denies the possibility of a positive metaphysical account of hallucination, it follows that *there is nothing in virtue of which* the hallucination is indiscriminable from a possible perceptual experience. The standard assumption which Martin is anxious to reject is that hallucination possesses, prior to our introspective

reflection upon it, some phenomenology that grounds its indiscriminability from some possible veridical perceptual experience. Instead, Martin contends that the indiscriminability is not reducible to the occurrence of any other mental properties - the hallucination has no *mental* or *psychological* properties grounding the indiscriminability.³⁸

It is worth noting Martin's wider metaphysical commitments in which he sets his negative epistemic account. He is expressly concerned to defend naïve realism while respecting what he calls *experiential naturalism*, namely the idea that sensory experiences are part of the natural realm of causes and effects (2004: 39). In the case of hallucination this means that whatever we appeal to in explaining the phenomenon must not include influences or constituents from beyond the natural realm (ibid.).

Martin is also particularly concerned to address the plausible possibility of 'proximally causally matching hallucinations', *viz.* that the very same kind of brain state might be involved in a perceptual experience and a (presumably subjectively indiscriminable) hallucination (2004: 53, 71).³⁹ It is crucial to understand why such causally matching hallucinations play a central role in Martin's theory. Plausibly, there could be hallucinations lacking the same proximal causes as possible perceptions – that is, cases in which the hallucinating subject's brain state is different from that involved in any possible perception. Such cases would still count as sensory experiences on Martin's account provided they exhibit the crucial indiscriminability property. They nonetheless pose a less direct challenge to naïve realism than causally matching hallucinations. Of course, as

³⁸ Importantly, the indiscriminability relation that holds between hallucination and veridical perception is not symmetrical. Martin does offer an overarching definition of visual experience in general – intended to embrace both hallucinatory and perceptual experiences – in terms of indiscriminability from possible veridical perceptions. However, although visual perceptions meet this criterion – being indiscriminable from themselves – this does not exhaust their nature, which in fact rests on their possessing something extra that is lacking from hallucinations, namely genuine phenomenology (Martin 2006: 376). The defining feature of hallucination is thus entirely parasitic on this extra feature of the corresponding perceptual experience.

³⁹ It is implicit in Martin's assertion of experiential naturalism that *perceptual* experiences too are caused, and he admits that "there does seem to be a causal dependency of our visual perceptions on the activity of parts of our brains, even if we do not yet know the full pattern of this dependency" (2004: 54). However, if we assume the relevant causes to be local, e.g. physical states of the subject, and also admit the possibility of 'proximally matching' hallucination then this is open to the objection that sameness of cause entails sameness of effect such that a perceptual experience and its proximally matching hallucination must involve the same kind of psychological state or event (Robinson 1994: 154 ff.). Martin argues in defence of naïve realism that perceived objects can be both causes and constituents of perceptual experiences, where something's being a constituent is, in this context, taken to be equivalent to its being a necessary condition on the occurrence of the experience (Martin 2004: 55-7; see also Martin 2006: 368). An alternative is to take perceptual experiences to be materially constituted by the whole physical chain of events linking object and subject, as per the simple metaphysical picture. It is doubtful then whether perceptual experience should be considered a state or event that is caused.

described earlier, any non-perceptual experience which seems to possess phenomenology brings into play the pressure to find a single, unified account of phenomenology regardless of what kind of experience possesses it. However, proximally causally matching hallucinations present an additional and more direct challenge, since here it seems that some event or state (in the subject's brain) is shared by both the hallucinatory and perceptual experiences. If that is so, and this event or state is sufficient for phenomenology in the hallucinatory experience, then it ought to be sufficient for (presumably matching) phenomenology in the perceptual experience, leaving nothing for the 'naïve realist aspects' – the direct awareness of worldly objects – to contribute phenomenology is it, the phenomenology yielded by the inner state or event common to both hallucination and perception would 'screen off' the awareness of worldly objects, leaving them no role in explaining the phenomenology in the perceptual case (Martin 2004: 62). It is these cases of causally matching hallucinations, vulnerable to this screening-off problem, that motivate and are the key target of Martin's negative epistemic account.

Causally non-matching hallucinations do not present such a direct challenge to the naïve realist. If the hallucinating subject's brain state is different from that involved in any possible perception, then the naïve realist might speculate that this distinctively non-perceptual brain state somehow 'generates' its own phenomenology. We could in such cases attribute the seeming phenomenology to something other than mere indiscriminability from possible perceptual experience (to awareness of sense data or qualia, for example), and to do so without directly raising the screening-off worry. This would yield a 'mixed' account of phenomenology, according to which the phenomenology of perceptual experiences is constituted by (awareness of) worldly objects, that of causally matching hallucinations is merely apparent, and that of non-matching hallucinations is constituted by (awareness of) something other than worldly objects, such as sense-data or qualia. Such a mixed view would avoid the screening-off worry, but would still face the objection that it presents a glaringly disunified account of something – phenomenology – that appears on reflection to have a common nature across the various categories of experience, perceptual and non-perceptual.

A second version of the no phenomenology view is offered by Fish (2009: 93 ff.). Fish's account, like Martin's, denies that hallucinations possess what he calls 'robust' phenomenology. A hallucination as of an F is, on this view, a state which causes the same cognitive acts (beliefs, desires, etc.) as would have been caused by the subject's veridically perceiving an F. These may include forming certain beliefs about the existence of an F before one, as well as the higher order belief that one is indeed veridically perceiving an F. Fish takes his account to be a development of Martin's theory, one which expands on what might naturally seem to be the implications of the indiscriminability criterion. *What it is* for a mental state to be indiscriminable from a possible perceptual experience is, on this view, for it to give rise to the same beliefs and judgments as would be produced by that perception.⁴⁰

Martin and Fish's accounts are, however, more different than this sketch might suggest. To understand the difference, it is helpful to look more closely at what, for Martin, is the central kind of case – the proximally causally matching hallucination. Here, we imagine a pair of cases, in both of which the subject's brain occupies the same kind of state. In the perceptual case this brain state is caused (distally) by the object perceived; in the hallucinatory case the brain state is caused (distally) in some other way, such as by an electrode-wielding neuroscientist. Now, although the perceiving and hallucinating subjects are in all relevant respects physically identical, the naïve realist assumes that they differ mentally in an important respect. Indeed, the naïve realist may think that they differ to the extent that the hallucinating subject is not enjoying a truly mental state at all. To see why, consider that the naïve realist is likely to think that the conscious mental aspects of perceptual experience are exhausted by its sensory phenomenology, and that this phenomenology is furthermore exhausted by awareness of the worldly objects perceived with their perceptible qualities. In the case of the proximally causally matching hallucination there is no awareness of worldly objects, so there is no phenomenology and no conscious mental aspects of the experience.

⁴⁰ As Fish acknowledges, we need to be careful when we talk of 'same cognitive effects' in this context. While we can assume that a proximally causally matching pair of hallucinatory and perceptual experiences would have the same *physical* effects, it is far from clear that their *cognitive* effects would be the same. A difference arises if we take these cognitive effects to be mental states (beliefs, for example), if we further assume those states to have their content essentially, and also allow that the content of those mental states is in at least some instances externally determined. A pertinent example would be perceptual demonstrative thoughts about objects – pertinent especially because a desire to account for the possibility of such thoughts is a key motivation for naïve realism in the first place. If such thoughts are externally determined – if their content includes reference to worldly particulars – then the corresponding thoughts arising from the proximally causally matching hallucination would lack the singular, object-dependent element of content, and so be different thoughts. The members of the perceptual/hallucinatory pair would therefore not give rise to all of the same cognitive effects despite their identical physical effects (Fish 2009: 94 fn. 13; see also Martin 2004: 63-4; Martin 2013: 42-3).

Thus, if the hallucination lacks any positive mental characteristics, as Martin insists, then there is indeed nothing positive to say about it in mental terms. Of course, in setting up the case we have acknowledged that we can say something about it in *physical* terms. The pressure to talk about it in psychological terms comes because the subject inevitably behaves as though she is enjoying a state with positive mental characteristics – a perceptual experience, to be precise. This similarity of behaviour is inevitable so long as we assume some form of determinism since, *ex hypothesi*, the subject is in the same physical state when hallucinating as she would be in if she were perceiving.

Now we can see the difference from Fish's account. Fish thinks that we *can* say something positive about hallucination in mental terms: we can say that something is a hallucination if it has the same cognitive effects as a possible perceptual experience, or at least some relevantly circumscribed subset of those cognitive effects. Fish in fact goes further, and says that a state is not hallucinatory *until and unless* those cognitive effects arise. In response to Siegel's objection that we can conceive of a subject having a hallucination but without manifesting some or all of the relevant cognitive effects (Siegel 2008: 217), Fish replies as follows:

It is the presence of the right kinds of cognitive effects that turns an otherwise unexceptional mental event into a hallucination. The hallucination itself has no special features at all that could enable us to identify it as a hallucination in the absence of the right kinds of effects. This means that, if cognitively sophisticated subjects do not believe they are having an experience of a certain kind, then they are simply not hallucinating.

Fish 2009: 104

The cognitive effects thus appear on Fish's view to be *constitutive* of the hallucinatory experience and not merely consequences of it. A hallucination is therefore a state that causes certain cognitive effects, and has those causes essentially. For Fish, the obtaining of the relevant cognitive effects (in the absence of phenomenology) is both necessary and sufficient for the occurrence of a hallucination. This is certainly not Martin's view – while he might accept that a hallucination has some or all of the same mental and/or physical effects as a proximally causally matching perceptual experience, he would not treat these as constitutive of the hallucination itself. By restricting himself to a negative epistemic criterion for hallucination, Martin does not need to claim that in order to hallucinate a subject must exercise any or all of a certain set of beliefs or judgments.

For Martin, the occurrence of certain cognitive effects is not sufficient for a state's being a hallucination, since it might be that these effects are caused by a state that is not hallucinatory in the required sense, a state in which the subject *could* know she was not perceiving if her powers of reflection were more acute. Nor should Martin think the occurrence of certain cognitive effects *necessary* for a state's being a hallucination even if, as noted above, we should expect certain beliefs to arise (chiefly, that she is perceiving some *F*). After all, Martin's is a modal claim; a subject is hallucinating so long as she *could not* know by reflection alone that she is not perceiving. It does not require that she *does in fact* come to form a belief to that effect; she may for example, pass out immediately after having the hallucination and before she has been able to form any relevant judgments.

The obvious objection to both Martin and Fish's accounts, for all their ingenuity, is simply that they fail to do justice to our intuition that hallucinations have genuine, robust phenomenology. In other words, they do not seem adequately to account for the felt reality of hallucination. Plainly, since both Martin and Fish base their account precisely on the denial that hallucinations have robust phenomenology, this objection is unlikely to move them much. What's more, both argue that they can account for the felt reality of hallucination without appealing to robust phenomenology. According to Fish, this is explained precisely as the outcome of the subject's undergoing the same cognitive effects as would result from an appropriate perceptual experience. That is, the hallucination's feeling real is nothing more or less than its being constituted by those cognitive effects – for it to feel real is just for the subject to believe that she is perceiving (2009: 97-8). The felt reality of the hallucination is thus parasitic on the felt reality of perceptual experience (2009: 109-10). Understanding the parasitic nature of this relationship avoids the possible objection that the cognitive effects, being common to both perception and matching hallucination, ought to suffice for the felt reality of the perceptual experience too. The felt reality of perceptual experience is granted by its robust phenomenology – its presentation of worldly objects and their qualities – and the cognitive effects in the case of hallucination impart a felt reality only by reference to that robust phenomenology. The hallucination

lacks robust phenomenology but the subject nonetheless thinks and acts as though it were present.⁴¹

Martin similarly argues that his negative epistemic account can explain the felt reality of hallucination – as he puts it, the indiscriminability property definitive of (conscious) hallucinatory experience can be treated as determining what it's like for the subject to have that experience (2006: 397).⁴² He rejects the suggestion that the negative epistemic account leaves hallucinatory experience as a failure to recognise the absence of phenomenology (as a case of "unknown absent qualia" as he puts it), instead arguing that we should treat it as a case of phenomenal or sensory consciousness understood in terms of the subject's 'subjective perspective' (Martin 2006: 396). Fish in fact says very much the same thing when defending his notion of 'felt reality' from Martin's objections: he says that felt reality is not to be construed as something a bit like phenomenology which accompanies certain cognitive states, but instead it is intended to capture merely the subject's (misleading) perspective on her own mental state. It is "intended to express how things are for the subject – how things are from the subject's point of view" (Fish 2013: 63). For Martin, this subjective perspective is one characterised negatively, as a limitation on what the subject can know about her own mental state, one which is nonetheless sufficient, according to Martin, for its seeming to the subject as though she were perceiving. For Fish, the subjective perspective can be characterised positively, as the subject's holding certain beliefs about her mental states, beliefs that she would have had were she perceiving.

⁴¹ Martin argues that Fish's account of the felt reality of hallucination fails because Fish offers no account of what it is in virtue of which hallucinations have the relevant cognitive effects (Martin 2013: 45). However, Martin's own appeal to sameness of proximal neural effects would suffice here: hallucinations have the same cognitive effects as possible perceptual experiences in virtue of involving the same kind of brain state (see Martin 2013: 46). Note that appeal to the involvement of a common brain state does not require a commitment to that brain state being a 'realizer' of either the perceptual or hallucinatory experience, contrary to what Martin states here (ibid.).

⁴² I have avoided using the phrase 'what it is like' in connection with Fish's theory of hallucination. In his (2013) he states that his definitional use of 'what it is like' is tied to the notion of phenomenology, which in his account is the assemblage of robust properties of experiences that determine their conscious character. This means, he suggests, that if there is something it is like for the subject to be in a certain state then that state has (robust) phenomenology (Fish 2009: 8-11; Fish 2013: 62-3). The issue is somewhat complicated, however, by his use of 'presentational properties/character' to denote the (assemblage of) properties "that are, *or at least seem to be*, presented to the subject of experience and thereby characterize *what it is like* to be in the experiential state" (2009: 12; my italics). It would therefore seem open to him to claim that there is something it is like to hallucinate since, although lacking phenomenology (there being no robust experiential properties), hallucinations nonetheless possess presentational character since they *seem* to present the subject with properties.

Whether or not we find Martin and Fish's explanations of the felt reality of hallucination convincing, we are left with the related problem that both Martin's and Fish's accounts rule out the most natural explanation of why a hallucinating subject comes to entertain the relevant beliefs. The natural explanation is to say that it is the phenomenology of the hallucination that induces the subject to believe that there is an Fbefore her. But on Fish's account the entertaining of the relevant beliefs is constitutive of the hallucination and so cannot be a consequence of it (see Pautz 2013: 30). On Martin's account, meanwhile, hallucination consists simply in the brute unknowability condition which is clearly incapable of explanatorily grounding any claims to knowledge. For these same reasons, it will not do to appeal to the felt reality of hallucination (to use Fish's term) to explain why the relevant beliefs arise, for the fact that it seems to the subject as if she is perceiving some F is *constituted* by the indiscriminability or the relevant cognitive effects (as the case may be), and cannot thereby explain them. Of course, given the nature of their accounts, neither Martin nor Fish feels compelled to provide an explanation of the beliefs consequent upon (or constitutive of) the hallucination – the hallucination has no substantial metaphysical nature that could ground either the beliefs or the explanation. However, this merely adds to the counter-intuitive consequences of denying phenomenology in these cases.

Before leaving the no phenomenology view, it is worth remarking on the scope of Martin and Fish's accounts of hallucination. With their appeal to indiscriminability and sameness of cognitive effects, Martin and Fish's accounts focus on hallucinations that seem to 'match' possible veridical perceptions; those that they call 'pure' or 'perfect' hallucinations. We might think that there are all manner of possible hallucinations that are nothing like genuine perceptual experiences. These might include stereotypically 'fantastic' hallucinations involving unworldly juxtapositions of elements (pink elephants, for example), as well as the sort of visual experiences sometimes induced by migraine, anoxia and certain drugs, often taking the form of geometric patterns appearing to 'float' in the visual field. At first sight, it might seem that such cases would not even count as sensory experiences given Martin's definition of sensory experiences (intended to cover both perceptual and hallucinatory cases) as any experience that a subject cannot know by reflection is not a veridical perception. (Lucid dreams would seem to offer a good example of a sensory experience that is knowable by reflection alone to be a non-perceptual experience – it is precisely definitive of *lucid* dreams that the subject knows she is dreaming, and this despite the vividness of the phenomenal qualities experienced.)

However, both Martin's and Fish's accounts arguably have the scope to accommodate these kinds of cases as hallucinations. In Martin's case, the matter hinges on the scope of the 'by reflection alone' clause in his definition of sensory experiences in general. As noted earlier, the clause is intended to allow for the possibility that a subject is hallucinating and knows that she is not perceiving on account of some background knowledge (perhaps she remembers taking the LSD or knows that the friend she is hallucinating is currently in another country). To put it another way, all that would be required to show that the subject could not know by reflection alone that she is not perceiving would be establish that there *could* be some perceptual experience, however contrived, that matched the hallucination in question. Even in the case of the hallucinated pink elephant or the migraine aura, we could imagine some carefully stage-managed scenario in which a perceptual experience could be made to match those experiences. Fish's account lends itself to a similar scope-broadening claim: for him, it hinges on the possibility of perceptual experiences that would elicit the same cognitive effects, and the possible contrivances just mentioned would do precisely this.⁴³

⁴³ In fact, Martin offers an alternative way in which we can accommodate within the general picture those hallucinations that might seem obviously hallucinatory on reflection. Those he mentions include hallucinations of impossible scenes (an Escher staircase), novel colours (supersaturated red) and, less dramatically, the sorts of partial hallucination in which hallucinated elements are experienced against a background scene that is genuinely perceived (2004: 80). His suggested explanation is that we could apply the indiscriminability criterion not to the experience as a whole but to discrete elements of the experience (2004: 81). To take an example, we might know by reflection alone that a hallucination of an Escher staircase is not a veridical perception, but we could segment the experience into elements that, although impossible when combined a certain way, are potentially veridical when taken in isolation. Two things are worth noting here. First, which cases are problematic in this respect depends partly on what scope we grant to the 'by reflection alone' and so whether we think that, e.g., recognising the impossibility of an Escher staircase relies on background knowledge. Second, and more importantly, the need for an appeal to experiential segmentation arguably arises because Martin arranges his theory around the contrast between hallucination and veridical perception, i.e. leaving illusory perception unaccounted for. If we assume that illusory perceptions can be fully explained by reference to how things are in the worldly scene before the subject, then there should be no reason for Martin not to broaden his central claim so that the concept of a sensory experience is general is of an experience that is indiscriminable from a possible perception, irrespective of whether that perception is veridical or illusory. This would allow us to explain the hallucination of an Escher staircase in terms of the experience's indiscriminability from a perception of a drawing of an Escher staircase or a carefully crafted model of the same, even if we consider these to be illusory perceptions. It is unclear what to say of the supersaturated red case, although it is not obviously explained either by Martin's appeal to segmenting experience into discrete elements, since it is unclear how we could separate a hallucinatory experience of novel colour into discrete 'elements'.
To conclude this section, we have seen that the no phenomenology view seeks to avoid the problems consequent upon any attempt to give a positive explanation of phenomenology in terms of awareness of something other than worldly objects. However, in so doing, its central claim runs counter to our intuitive commitment to the notion that hallucinations can have a substantial felt reality grounded in genuine awareness of properties. In the next section I will describe an alternative account of hallucination, based on the simple metaphysical picture, which promises to avoid the argument from hallucination while acknowledging that hallucinations have robust phenomenology.

4.6 A simple metaphysical picture of hallucination

On the face of it, the simple metaphysical picture might seem worse than useless for explaining hallucination. Its explanation for the object-involving nature of perceptual experience works by taking what is standardly treated as a *cause* of the experience and treating it instead as a *constituent* of the experience – by taking the experience to be, not a mental and/or physical event within the subject, but the *process* that culminates in some physical event within the subject. But in the case of hallucination we assume that this involves an experience that expressly *lacks* a real worldly object. It would thus appear that the simple metaphysical picture is useless to serve the explanatory purpose it serves in respect of perceptual experience, where it identifies the elements of phenomenology with real worldly constituents of the perceptual process.

Invoking the simple metaphysical picture might indeed seem to make matters worse if we further assume that the brain state involved in hallucination must have *some* causal antecedents – psychedelic drugs or a devious neuroscientist wielding electrodes, perhaps. Applying the simple metaphysical picture to such cases generates the puzzle why the hallucinatory experience is not ostensibly *of* those causal antecedents (assuming it isn't) – why, in other words, the drugs, the scientist or the electrodes do not feature as elements of the hallucinatory phenomenology.

I will argue, however, that the simple metaphysical picture can underpin a positive account of hallucinatory phenomenology after all. It does so by showing how we might reject the first assumption, namely that hallucination lacks a real worldly object. The core idea is that a hallucinatory experience as of an F might in fact be constituted by a process

in which an *F* is a causal determinant of the subject's hallucinatory brain state.⁴⁴ Of course, this is much the same sort of claim that the simple metaphysical picture makes for perceptual experience – that a perceptual experience of an *F* is constituted by the process in which that *F* visually affects the subject. The difference is not that the hallucinatory experience *lacks* the *F* as a constituent, but that the process linking the *F* and the subject is, in some important respect, atypical.

Consider, for example, a typical case of visual experience – of seeing a cat, say. The subject's seeing the cat involves a process in which the cat causally influences the subject (most pertinently her brain state) via appropriate causal intermediaries (reflected light, retinal states, optic nerve firing, etc.). This is not yet to say anything about the nature of the experience itself. According to the causal theory of perception, this experience is the final upshot of the causal process described, leaving a question mark over the relationship between the phenomenology of the experience and the object perceived (the cat) which is merely a cause of the experience' is nothing more or less than the *experiencing*, which is an active relation of awareness holding between subject and object and which is describable in physical terms as a process of a certain sort. This allows us to envisage a straightforward relationship between the phenomenology of perceiving and the worldly objects perceived: the former is simply a matter of how those objects' qualities are arrayed within the subject's field of view.

Now consider a case in which the subject hallucinates a cat. For simplicity's sake, let us suppose that the subject hallucinates her pet cat Oscar. What does the phenomenology consist of in this case? I suggest that, as in the perceptual case, the relevant element of the hallucination's phenomenology is nothing other than Oscar himself. Of course Oscar is not seen, and we shall assume that Oscar is not currently located within the subject's potential field of view. In fact, let us assume that Oscar is in fact dead and buried. It might seem therefore that our subject cannot be standing in any kind of relation to Oscar now. However, this would be too hasty: it is perfectly possible that the subject is currently standing in some (causal and/or psychological) relationship to Oscar as he was at some

⁴⁴ The proposal here is structurally identical to that made by Manzotti in what he calls his 'process-oriented view of perception' (2008: 187-8), although the upshot is metaphysically quite different. As I explained in chapter 1, Manzotti claims that phenomenal qualities are constituted by perceptual processes, whereas I claim that phenomenal qualities are, at least in many 'good' cases of perception, where they seem to be, inhering in worldly objects.

time in the past when he was alive. That is precisely what I suggest - just as in the perceptual case, the subject's physical state, and in particular her brain state, is appropriately (i.e. visually) a causal consequence of Oscar, albeit as he was at some time significantly past rather than in the present. How can this be? Since Oscar is dead and gone, and since his effects in structuring the distribution of reflected light in the subject's environment have long since dissipated, we must look inside the subject's physiology for any residual effects of Oscar. From here the story plausibly looks something like this: the structure of the subject's brain retains the effects of Oscar as, *inter alia*, some visual memory trace. The details of how this might work physiologically are unimportant, provided that the subject's brain has some physical properties that would be lacking in the counterfactual scenario in which she had never seen Oscar, and which embody information about Oscar's visible qualities. Next we suppose that, for some reason, this memory trace becomes 'activated' in such a way that it interacts causally with other brain states of the subject that are involved in reflective awareness and belief formation. Typically this sort of interaction would amount to the subject's enjoying a visual memory of Oscar and recognising her mental state as such. However, for whatever reason, on the occasion in question the subject forms the erroneous belief that Oscar is not merely recalled but is actually seen, perhaps because the activation of the memory trace results in an unusually vivid visual recollection of Oscar. The subject is now hallucinating; it seems to her that she is seeing Oscar.

It is implicit here that at least some kinds of memory – specifically, episodic or 'imagistic' memories – are quasi-sensory experiences with (perhaps vague or indistinct) phenomenology, and that this phenomenology is to be explained via an extension of the simple metaphysical picture. Extending the simple metaphysical picture amounts to claiming that there is nothing more to (perceptually or 'imagistically') remembering something than one's standing in a certain relation of (delayed or deviated) awareness to that thing, a relation fully describable in physical terms as a process involving, *inter alia*, the object's sensorily affecting the subject. As I have explained it, this effect will include the establishment of appropriate 'memory traces', and the 'reactivation' of these memory traces' in such a way as to restore the subject's conscious access to some of that object's perceptible qualities. Weir and Manzotti separately propose just this account of experiential memory as part of their 'ultra-realist' and 'process-oriented' theories of perception, respectively. Manzotti explicitly describes memory precisely as 'delayed awareness' effected via persistent modifications to the brain resulting from earlier perceptual encounters, i.e. memory traces (Manzotti 2006: 26-7). In Weir's terms, the reactivation of memory traces in the brain "bind[s] earlier objects and properties once more into situations" (Weir 2004: 120). Both Weir and Manzotti exploit this account of memory in order to explain hallucination

Note that the appeal to memory in explaining hallucination is plausible independently of applying the simple metaphysical picture. The thought is just that in all actual cases of hallucination, the phenomenology depends on and borrows from that of our past experiences, if not wholesale then at least in terms of its constituent elements. Importantly, as both Weir and Manzotti point out, the possibility that hallucinatory phenomenology is sometimes 'assembled' out of elements of disparate perceptual experiences allows the account to explain not only memory-based 'snapshots' of previously perceived scenes but also the more fantastic or bizarre hallucinations that we might think of as stereotypical (see Weir 2004: 122-3; Manzotti 2008: 188). To use Manzotti's example, we might hallucinate a pink elephant despite never having seen such a thing because we have nonetheless seen diverse pink things and elephants. The subject's having the hallucination is due to the convergence of multiple processes originating variously in pink things and elephants, such that when the respective memory traces are reactivated their objects seem to coexist spatially, or rather the recollected quality of pinkness seems to qualify the elephant (Manzotti 2008: 188).⁴⁵ Or, as Weir puts it, the activation of the subject's memory trace binds diverse objects and qualities previously perceived into novel situations in which they may appear juxtaposed in ways never previously encountered (Weir 2004: 122-3).

The notion that the more fantastic hallucinations are 'assembled' by way of the joint re-presentation of diverse objects previously perceived might seem unattractively complicated. However, it is worth recognising that a similar complexity is likely to attend any alternative account of hallucination, even one rejecting the natural view in favour of, say, intentionalism. Suppose, for example, that one were to hold that hallucination results

⁴⁵ Note that, whereas Weir's account of memory and hallucination is set within the context of a theory consistent with the natural view, Manzotti's account holds that phenomenal qualities are themselves identical to the whole perceptual process involving subject and object, and so are not taken specifically to qualify the worldly object perceived.

from the inappropriate activation of 'perceptual capacities' that normally function to single out and recognise objects of a certain sort (Schellenberg 2013). Presumably, then, to explain a hallucination of a pink, winged elephant we must invoke the simultaneous activation of perceptual capacities for, variously, pinkness, wings and elephants. Not only that, but we must explain how the pinkness, wings and elephant components come to appear as combined in the appropriate fashion; in other words, mere simultaneity of activation of each perceptual capacity is not enough, else the subject might equally hallucinate an ordinary-looking elephant next to some unattached wings and a 'floating' patch of pinkness. In my own explanation of hallucination I have not offered a specific explanation of this apparent conjoinedness – my point is merely that it no more (and perhaps no less) faces this challenge than rival theories of hallucination.

I should acknowledge that, even if it seems plausible to appeal to memory in explaining hallucination, one might find it harder to swallow the claim that a past event or state of an object could be a constituent of an experience had now. It is worth noting, however, that this implication already falls from the simple metaphysical picture even as it is applied to perceptual experience. The perceptual process is of course a temporally extended chain of events such that what is perceived is always the object as it was prior to the moment of successful perception, assuming that the event of successfully perceiving is no earlier than the onset of the brain event caused by the object perceived.⁴⁶ In most ordinary situations this lapse of time between the event perceived and the act of perceiving is negligible, but there are perfectly everyday (if not quite mundane) perceptual experiences that involve a very substantial time gap. For example, to see the Sun is to see it as it was roughly eight minutes ago, given that this is the time it takes for the Sun's light to reach the Earth. We can see some other stars as they were thousands of years ago, so that we can, as it were, look back to a time long before we were born. This suggestion that a visual experience we are having now could take as a constituent some event that happened thousands of years ago is likely to seem ludicrous, although this reaction follows from taking 'experiences' to be ('mental') states of the subject, a notion that I have rejected. Suffice it to say for now that the awareness of past objects or events is a general

⁴⁶ In this respect, we often use 'perceiving' in its more specific cognates (e.g. 'seeing') to denote an achievement on the part of the subject, something that is datable roughly to the onset of the relevant effect on the subject's brain. *Seeing* in this sense is not to be equated with the visual experience, which takes the worldly object as a constituent and is, according to the simple metaphysical picture, not the outcome of the perceptual process but the process itself.

consequence of the simple metaphysical picture and not one required only by its account of hallucination as I present it here.

We should admit, all the same, that the nature of the time delay in cases of hallucination is somewhat different to that involved in, for example, seeing distant celestial objects. In the case of hallucination I have suggested that the delay is not to do with transmission of the perceptual medium (light, sound, etc.) across large distances, but is instead a feature of the memory-based mechanism proposed. Accordingly, the delay results from 'storage' of what was originally a perceptual effect in the form of a persistent change in the structure and/or function of the brain (the 'laying down of a memory trace', as we might put it). The causal relation holding between object and subject in hallucination therefore has this added complication or convolution when compared with the mere temporally drawn-out nature of the causal relation in cases of stargazing.

Even if we grant the claim that hallucinations are not unique or special in involving past events or objects, we might still wonder what it is about certain time-delay cases (seeing the Andromeda galaxy) that makes them count as instances of perception, while others (having a vision of Oscar) are treated as hallucination. The first thing to say here is that we should not expect necessary and sufficient conditions for distinguishing perceptual from hallucinatory experiences, any more than we should for distinguishing memory and hallucination. A key purpose of applying the simple metaphysical picture to hallucinatory experiences is precisely to show that they occupy some vaguely-bounded region of a continuum also occupied by perceptions and episodic memories, thus allowing us to offer a unified explanation for the phenomenology possessed by all such experiences.

Nonetheless, we do have some more or less common-sense basis for our ordinary judgments about whether an experience is perceptual or hallucinatory, just as we do for the analogous judgment about hallucination versus memory. For one thing, there is a question of *directness*: we expect perception to give us relatively direct awareness of objects, e.g. by the uninterrupted transmission of light from object to eye. Note that we shouldn't hesitate to say that we see the Andromeda galaxy directly when we look straight at it, even though the light has travelled for two-and-a-half million years to reach us, such that we are seeing the galaxy as it was two-and-a-half million years ago. After all, there is no *more* direct way to see it (although admittedly our ordinary attachment to the immediacy of seeing

might be strained here). The more interrupted or convoluted this transmission becomes, the less inclined we will be to call a case one of perception. Where the convolution is merely a matter of our seeing something via a reflection – the light bouncing off a mirror on its way between object and eye – we may be happy still to treat this as a case of perception. Less so once the transmission is mediated by television cameras, radio signals and TV sets, and less still once transmission is delayed through the storage of images on video recorders or indeed through the laying down of 'memory traces' in the brain as in the mechanism proposed for hallucination.

With cases like television viewing, another feature besides directness comes into play, namely *fidelity*. By 'fidelity' I mean the degree to which an experience accurately and completely presents the perceptually available elements of a scene. Some experiences will be treated as poor cases of perception because of some 'loss of signal', as when our vision is blurry immediately upon waking. More relevant to the distinction between perception and hallucination would be those cases of, on the one hand, partial hallucination, and on the other, fantastic hallucination. In the former case, the phenomenology includes more than the elements (objects and qualities) that are actually present within the environment in front of the subject. In the latter case, these extraneous elements are also juxtaposed with each other in ways that the subject has never seen, or indeed ways in which they may never have been juxtaposed at all, perceived or not.

A third feature relevant to the perception/hallucination distinction is the experience's contribution to a subject's powers to *track* worldly objects: when an object is straightforwardly seen the subject is typically in a position to track the object visually. In other words, changes in the visible qualities of the object (including its location as well as features like colour, reflectance, etc.) will, for as long as the subject is seen, correspond with changes in the subject's experience. In the hallucinatory case, meanwhile, we needn't expect changes in the subject's experience to correspond to actual changes in the object hallucinated, even if the hallucination were to take a single particular as its object. It could be, of course, that some hallucination takes the form of a faithful recollection of some previously perceived event, unfolding in the same sequence and with the same lapse of time as when originally perceived. In such a case, both the fidelity and tracking criteria are met, but the directness is not – the event was tracked more closely in the original perceptual experience.

Paradigmatic cases of perceiving, then, would be experiences in which our awareness of an object and its qualities is as direct as possible, and as complete and accurate as our perspective permits, and which enable us to track changes in the object. Paradigmatic cases of hallucinating would include those that meet none of these criteria. Many cases will fall in the middle and may elicit uncertain judgments about their classification as perception or hallucination. To repeat, it is not claimed that we can find here any necessary or sufficient conditions for either perception or hallucination, and by which we might sharply distinguish them. No such sharp distinction is to be expected since the same metaphysical underpinning is proposed to explain both kinds of experience.

We can offer similar observations to underscore our ordinary distinction between hallucination and (episodic) memory, as between hallucination and perception. What marks the difference between episodic memory and hallucination might include the vividness of the experience, the degree to which it is felt to arise voluntarily, and the degree to which it is mistaken for a perceptual experience. Again, we needn't look among these dimensions for some necessary or sufficient conditions for being a memory rather than a hallucination or *vice versa*. Since both admit of the same metaphysical explanation we can be content to allow that experiences might exist on a continuum between full-blown hallucination (vivid, unbidden and confused for reality) and obvious memory (indistinct, willed and recognised as recollection). For example, although forming the belief that one is perceiving is a typical consequence of many hallucinations, it is far from necessary; a subject can hallucinate knowingly, i.e. with insight into the non-perceptual nature of the phenomenology of her experience.

4.7 Dreams and other quasi-hallucinatory phenomena

It should be obvious that the simple metaphysical picture might form the basis for a similar explanation other experiential phenomena that we might loosely describe as hallucinatory, such as dreams, sensory imagination, after-images, and so on (see Manzotti 2008: 185 ff.). The examples I cited when discussing the no phenomenology view included phenomena such as migraine auras and the sort of geometrical or fantastic imagery arising through anoxia and certain forms of intoxication. Here the hallucination is often not one in which the subject seems to be aware of concrete objects, but one in which

there seem to be 'free-floating' patterns of colours that may not even appear to reside in the environment perceived. Such cases might be described as involving visual sensation rather than perception.

Since such hallucinations do not appear to acquaint us with a familiar object, or even a more or less simple 'patchwork' of (bits of) familiar objects (as with Manzotti's pink elephant), a simple appeal to memory-mediated awareness of previously perceived objects might seem less promising. Nonetheless, the same form of explanation does still suffice. Imagine a case in which subject *S* has a hallucination in which a circular grid of red lines occupies the centre of her field of vision. Let us further suppose that this hallucination is associated with some pattern of neural activity in the visual cortex of her brain (cf. ffytche et al. 1998). Although this pattern of neural activity, taken as a whole, has presumably never been associated with a perceptual experience, and is therefore not tantamount to the reactivation of a specific memory trace, we might plausible suppose that the structure and activity of each of the individual neurons involved, their connections to one another, and the structure and activity of the relevant brain region as a whole, have been influenced causally by the innumerable visual experiences that the subject has enjoyed and which have involved that brain region. In other words, the structure and functioning of that region is a causal consequence of both the genetically determined development of the brain since the subject's conception and the diverse 'traces' left behind by its visual interactions with the subject's environment over its lifetime so far. Then, when some local cause (oxygen deprivation or a psychedelic drug, say) sparks the relevant pattern of visual cortical activity, that activity is also an effect of diverse objects previously encountered perceptually. We might speculate that, in the example proposed, a feature common to many or most of these previously seen objects is their possessing the colour red; the activated neurons might be ones whose pattern of activity in hallucination is what it is as a result of their previous activation by the redness of perceived objects. Thus, when activated in hallucination, their precise form of activation is a causal product not just of whatever proximal causes brought about the hallucinatory brain state but also of more distal causes, crucially including the red objects previously perceived. It is the 'reawakening' of the visual-causal link between the subject and these red objects that accounts for the red qualities of awareness in the hallucination.⁴⁷

⁴⁷ It is a further question how many, and which, previously seen red objects one is thereby aware of. If we think of the awareness relation in causal terms it is easier to see how there might be a matter of fact here,

We can approach the problem from a different direction if we return to the contention that migraine auras, after-images and the like involve visual sensation rather than perception. Here, one line of thought antagonistic to the natural view might be that if we can treat such cases as involving sensation rather than perception and further acknowledge that these present genuine phenomenal qualities, then we ought to treat paradigmatic perceptions as also composed ultimately of sensations, albeit sensations that we 'project' onto the world or construe as worldly things and qualities. The proponent of the natural view should here turn the argument on its head, however. We can imagine a situation where someone blind from birth has his sight restored, and in which this formerly blind subject is at first unable to make sense of the shifting arrangement of visually-apprehended qualities now presented to him. We could of course take this as introspective evidence for the claim that perceptions are constructed from sensations. However, we could equally interpret the situation as one in which the subject is simply unable to correctly 'place' the objects or qualities of which he is now aware.

What I suggest is happening in cases of migraine auras, after-images and the like is that subjects are likewise unable accurately to locate the objects or qualities of which they are aware. However, in such cases the inability reflects no inexperience or lack of skill on the part of the subjects, as it does in the case of the formerly blind subject attending to the scene before him. In after-images (or, equally, migraine auras or geometric hallucinations) we are given no clues as to the location of the qualities of which we are aware. In normal perception, the worldly qualities occupy a more or less stable location relative to the subject and to each other, and – at least for things close at hand – share the same space and time. Some hallucinations too might give us clues regarding the worldly spatial and temporal location of the hallucinated qualities, as when the hallucination is effectively a highly vivid memory image or 'flashback' to some specific event. Migraine auras and the like are limit cases in which the arrangement of qualities is so thoroughly 'shuffled' that reconstruction or recognition of their original spatial or temporal location is all but impossible.

assuming some principled means of establishing the causal influence of any previously seen red object on the form of the hallucinatory brain state.

After-images comprise an especially interesting class of quasi-hallucinatory phenomena, whose explanation via the simple metaphysical picture is a little more complicated than those discussed so far. After-images result from the visual system's adaptation to prolonged exposure to a single colour stimulus.⁴⁸ They are explained as arising mainly as a consequence of the opponent structure of visual processing, allied to the phenomenon of 'neural rebound', whereby nerve cells caused by a stimulus to enter one state (excitatory or inhibitory) will 'rebound' into the opposite state when the stimulus is withdrawn (De Valois & De Valois 1997: 127). Accordingly, a neuron that responds ultimately to a red stimulus by entering an excitatory state will rebound into an inhibitory state when the red stimulus is withdrawn. As it happens, it is a feature of the opponent structure of the visual system that this inhibitory state is the neuron's normal response to a green stimulus (ibid.). The upshot, then, is that prolonged viewing of a red object will, when abruptly terminated, cause the subject's brain to enter a state that is otherwise caused by the viewing of green objects. According to the account I have proposed, consistently with the natural view, this reaction will not be sufficient for the appearance of green enjoyed by the subject. For this the subject must be aware of some *actual* instances of greenness. To deliver that, we must suppose that the subject's 'rebound' brain state is not only an effect of the (proximal)) red object but also an effect of previously seen green objects. And that will require an explanation like the generic one offered above, whereby the rebound state of the relevant brain region bears some 'trace' of previous causal interactions between that brain region and green objects.

Admittedly, this explanation leaves much unsaid about how or why traces should be left. However, this sort of explanation of after-images appears to be what is required if we are to retain the natural view along with naïve realism's explanatory virtue; that is, if we are to explain the phenomenology of after-images by appeal to awareness of actual instances of the relevant colours. As we will see in chapter 6, there are good reasons to doubt that this explanatory virtue, and naïve realism more generally, can hold in the face of some other phenomena of colour perception. In chapter 7 I will consider the implications of this for the natural view.

⁴⁸ I am concerned here with *negative* after-images, i.e. those of a colour complementary to the original stimulus. *Positive* after-images, i.e. of the same colour as the stimulus, are also possible, but are elicited by different means (De Valois & De Valois 1997: 126).

In the meantime, we might press the simple metaphysical picture into service in explaining a puzzling form of illusion that I mentioned in chapter 3.5, namely motioninduced colour, as in Benham's Disk. In that chapter I promised to offer an explanation borrowing the account of hallucination. The explanation here might be related to the one offered above for simple geometric hallucination. So, we might imagine that the rapidly rotating or flickering black and white pattern elicits, by some physiological quirk of the visual system, certain neural effects that are not normally induced by black and white stimuli but which in fact are more typically induced by things of the illusory colour (red and green, say). When the disk spins fast enough, then, a certain effect is brought about which is constrained in its form by neural modifications brought about by previous visual encounters with red and green things. The effect in question therefore reawakens the visual-causal links to those previously perceived red and green things and so brings their redness and greenness into the subject's awareness.

This account raises an interesting question. Since the neural activation is *now* the effect of a spinning or flickering black and white object, we might wonder whether this object is 'in competition' with the more distal causes (red and green things) for a dominant role in phenomenology. If a subject were granted innumerable Benham's Disk viewings but deprived of coloured objects, would the neural effect come to be constrained causally more by the (recent) encounters with the disk and less by the increasingly historical encounters with red and green things? If so, we ought to expect the illusory colour to fade from the disk experience. Certainly, if the simple metaphysical picture explanation is correct, we should expect that a subject with normal vision but deprived of colour experience birth would not be aware of any chromatic colour when exposed to Benham's Disk. We might in fact expect that such a subject would, if exposed to something red or green for the first time, have an experience whose phenomenology was in part constituted by previously perceived spinning or flickering black and white patterns, in a reversal of the normal illusion.⁴⁹

⁴⁹ One alternative explanation of the Benham's Disk illusion might be to argue that it is not an illusion at all, and that the disk *is* coloured when spun fast enough. We are likely to reject this on the empirical grounds that we find no difference in the physical surface properties (including spectral reflectance) when the disk is spun, and therefore we would lose any explanatory link between the measurable properties of the object and its observable colour. However, as I will explain in chapter 6, this explanatory link is one that is loosened by consideration of certain other features of colour and colour perception.

4.8 Missing and impossible colours

The account I have proposed in this chapter seeks to explain what it's like to hallucinate by appeal to the awareness of previously perceived objects (their 're-presentation', as I put it). Now, however, the account faces a problem with the possibility of one's hallucinating, imagining or illusorily perceiving qualities that one has never perceived before – such as Hume's missing shade of blue (Hume 2007: 14-15) – or qualities that are in fact not instantiated anywhere, such as 'supersaturated red' (Johnston 2004: 141-2) and impossible-sounding binary hues like reddish-green and blueish-yellow (Crane & Piantanida 1983; Billock & Tsou 2010).

Usefully, the explanation of the subject's awareness admits of a different explanation in each case. One can be understood as a case of illusion ('impossible' binary hues), another as a case of partial hallucination (supersaturated red), and the third as a more complex act of perceptual imagination (the missing shade of blue). We can take these in turn.

The awareness of 'impossible' binary hues can be explained as a form of illusion, namely as misleading awareness of qualities actually instantiated before the subject. Subjects can be made to see reddish-green and yellowish-blue by presenting subjects with adjacent vertical stripes of opposing colours (red and green, or blue and yellow) and by stabilising the boundary between the stripes on the subjects' retinas by means of eyetracking devices (Crane & Piantanida 1983; Billock & Tsou 2010). In the earlier experiments, subjects variously reported a single uniform field of colour, a granular array of small patches of each distinct colour, or larger 'islands' of one colour against a background of the other (Crane & Piantanida 1983: 1079). Later researchers discovered that the patchy appearance of grains or islands can be eliminated by using adjacent colours of equal luminance (Billock & Tsou 2010). This leaves the appearance of a uniform field of colour that is neither solely one colour nor another, which is the reddish green or yellowish blue of interest. Subjects vary, nonetheless, in their description of this colour field. Some find it difficult to describe, while recognising that it is coloured in some way. Most, however, describe them as simultaneously red and green or blue and yellow (Crane & Piantanida 1983: 1079). There remains some variation in how this simultaneity appears to manifest: in some cases as a gradient from, say, pure red at one side to pure green at the other, and in between a continuous range of greenish reds through reddish greens; others

see red and green fields as if superimposed on one another but at different depths; while others see just a uniform field of reddish green (Billock & Tsou 2010). It is clear that in at least some of these cases the subjects take themselves to be aware of both green and red at the same time. In other cases, this is less clear, and the authors seem to suggest that reddish green can sometimes appear as a single, uniform and unfamiliar colour and not just an amalgam of two familiar colours, perhaps in the way that purple can be understood as a mixture of blue and red without appearing simultaneously blue *and* red.

Regardless of which appearance results, Billock & Tsou conclude that although the subjects might be said to perceive 'new' colours, these remain "compounds of familiar colors" (2010). This would suggest that the outcome is consistent with the natural view that the phenomenology is wholly explicable as awareness of qualities instantiated in the subject's environment. Certainly, it is crucial to recognise that what is going on actual cases of seeing 'reddish-green' is genuine perception, albeit illusory perception. There is therefore no doubt that subjects are in fact simultaneously seeing red and green as if they were occupying the same region of the field of view (but perhaps as if they were at different depths). That this should look strange is no surprise once we acknowledge that seeing red and green as co-located in this way, but without them interfering with one another, is normally impossible. If all awareness of reddish green is illusory it is because red and green are never co-instantiated in worldly objects, but this does not preclude our being aware of red and green *as if* they were co-located.

Awareness of another 'impossible' colour – supersaturated red – admits of a different explanation, one which appeals not just to awareness of qualities currently instantiated before the subject but also to awareness of qualities previously perceived. This case is therefore better classed as a partial hallucination than an illusion. Supersaturated red is an unnaturally intense or vivid shade of red which subjects appear to see when they look at an ordinary red object after first looking at a bright source of pure green light. In effect, the subjects experience a red after-image 'superimposed' on the red object, and the appearance of unnaturally vivid, 'supersaturated' red is the result (Johnston 2004: 141-2). Again, the supposed problem for the natural view is that supersaturated red is a colour that is never in fact instantiated in the world. As such, it might seem to elude the explanation offered for hallucination, *viz.* by appeal to actual qualities previously perceived. However, here again we can offer an explanation of this sort after all. We might explain awareness of

supersaturated red as resulting from simultaneous awareness of multiple instances of redness as if they were co-located or superimposed. Given the mechanism of inducing the experience, this is expressly just what is going on, at least if we accept that enjoying a red after-image involves one's being aware of redness. I explained the latter as awareness of previously perceived instances of redness. As such, the actual red objects of which the subject is aware, in having the experience of supersaturated red, are not co-located in time and need not be co-located in space. One of the red objects is currently before the subject, but the other red objects – those contributing constitutively to the after-image – are objects seen at different times and most likely in different places.

Although we might thus explain supersaturated red in terms of the simultaneous awareness of different instances of redness, we should acknowledge that this is quite a peculiar form of simultaneous awareness of redness. Supersaturated red is not typically seen when, for example, we mix two lots of red paint, or view a red object through red lenses, or cross our eyes so that two red patches, each seen with a different eye, appear to overlap in our field of vision. Rather, this is a peculiar situation in which we are, as it were with each eye, simultaneously aware of two separate instances of redness as if co-located, a sort of 'double awareness' that is impossible via genuine perception. The result is not a novel *hue* but a familiar hue manifesting abnormal saturation or intensity.

The crucial question, ultimately, is not whether novel colours like reddish green or supersaturated red can really be instantiated in worldly objects but whether the appearance of those colours can be explained solely by reference to worldly objects and their perceptible qualities. The simple metaphysical picture offers us a way to answer this positively.

We might similarly offer a positive explanation of the possibility, entertained by Hume, that a subject might imagine a colour she has never seen before. Here Hume imagines a subject who has encountered all manner of colours, including many shades of blue, but who has never encountered one particular shade of blue. Hume considers it plausible to assume that the subject will nonetheless be able to *imagine* this 'missing' shade (Hume 2007: 14-15). Here, we should understand 'imagine' in terms of Hume's notions of 'impressions' and 'ideas', and so as involving a kind of awareness of the missing shade of blue that is 'less lively' than perception of an actual instance of that shade

(ibid.: 17-18). If we are to accommodate this possibility consistently with the natural view we might explain it via the simple metaphysical picture more or less along the lines of imagination more generally, viz. as some sort of 'rearrangement' of qualities previously perceived. In this case, the relevant qualities will be other shades of blue, as well as other colours depending on just where the 'missing' shade lies along a spectrum of possible colours. To see how this works, we can reflect on the observation that the various colours can be understood as either 'pure' ('unique') or 'mixed' ('binary'). The former include those colours – certain shades of red, green, blue and yellow – which do not appear to have any hint of each other. The mixed or binary colours, conversely, are those in which there appears to be some discernible hint of more than one colour, such as a bluish-green (see e.g. Allen 2010a, Tye 2006). On this model, the imagining of Hume's missing shade of blue can be understood as the 'remixing' of pure hues previously perceived, just as imagining (or hallucinating) a pink elephant involves the rearrangement of elements previously perceived separately. If there is a difference between the missing shade of blue and the pink elephant, it is that the latter seems more obviously to involve a simple juxtaposition of previously disparate elements. The missing shade of blue, meanwhile, is a uniform and novel colour rather than a juxtaposition of, say, yellows and blues previously seen. The account therefore relies on the assumption that we really can understand the missing shade of blue as being composed of two familiar colours, rather than being irreducibly a novel colour in its own right. It is not obvious how we could settle this definitively. Nor, however, is it obvious that Hume is right to think that a subject *could* imagine a colour never before seen.

4.9 Which constituents of the perceptual process are perceived?

Having seen the view spelled out with a variety of examples, we can now revisit my acknowledgement at the start of this chapter that, as a means of explaining hallucination, the simple metaphysical picture might appear useless, if not *worse* than useless. The mere uselessness was supposed to follow from the clash between the claim (consistent with the simple metaphysical picture) that hallucinations are object-involving experiences in essentially the same way as perceptual experiences, and the standard characterisation of hallucinations as object-less experiences. With the simple metaphysical picture-based account laid out we can now see that this *prima facie* objection is without merit – an

object-involving account of hallucination can be given while retaining most if not all of the ordinary conceptual distinction between hallucination and perception.

This should help to defuse the further worry that the simple metaphysical picture might be *worse* than useless as an explanation of hallucination. The concern here arose from the thought that a hallucinatory experience, even if it lacks a worldly object, nonetheless might be presumed to have *some* worldly cause, such as a psychedelic drug or a devious neuroscientist with his carefully placed electrodes. If these are the causes of the brain state involved in hallucination, why do they not feature in the phenomenology of the experience? With the recognition that the relevant causes can include just the sorts of worldly objects that seem to feature in the phenomenology of the hallucinatory experience, the sting is taken out of this objection to some extent, since we now have available a straightforward explanation for the phenomenology in terms of the presentation of worldly objects and their qualities. As a result, the simple metaphysical picture-based account of hallucination can lay claim to what I have called the explanatory virtue of naïve realism:

<u>The explanatory virtue</u>: naïve realism explains what it's like to perceive an object solely by reference to objects and qualities in the subject's environment as they are arranged to the subject's perspective.

This does not entirely neutralise the objection, however. Lurking within it is a worry that applies to the simple metaphysical picture's explanation of experience more generally, whether hallucinatory or perceptual. The worry is simply that identifying an experience, of whatever sort, with a physical process does not constrain which constituents of the process should feature as elements of phenomenology. Why should a process in which an apple visually affects a subject's brain amount to an experience whose phenomenology is composed of the apple? Why should the phenomenology not feature any of the other constituents of the process – the mediating light, the optic nerve or the brain state itself? Furthermore, any such process is open-ended historically, so we might wonder why the phenomenology doesn't feature elements of the causal process *preceding* the apple – the sun that shines on it, the tree from which the apple was plucked, the seed from which that tree grew, and so on. In the case of an apple hallucination that is instigated by the neuroscientist's electrode, the simple metaphysical picture-based account will claim that the electrode caused activation of a memory trace of an apple, thereby bringing into being

a new causal process mediating between the historical apple and the subject's current brain states in such a way that the subject is now able to reflect on the visual appearance of that apple. But why should the subject not now be able to reflect also upon the visual appearance of the electrode?

It should be clear that this is a puzzle that arises only from the simple metaphysical picture and the way in which it presents the awareness relation from an outside perspective, as if we were examining the physical goings-on involved in someone else's perceptual act. The natural view itself faces no such puzzle. It is not a claim about how an object gets to be 'in' experience – how it gets to be the thing perceived. The natural view, as a claim about phenomenology, is a claim about experience from the 'inside', i.e. about how things seem from a first-person perspective. From this perspective, the object is already there in our field of view, and the natural view is not concerned with how it gets there (simply by being seen) but with what it is that occupies our field of view (possesses phenomenal character), namely just the object itself. What the simple metaphysical picture offers is a way to 'reconstruct' how it is that we become aware of objects. If we were concerned only with veridical perception, we mightn't feel the same need for such a reconstruction, since the relationship between ourselves and the objects we perceive is quite transparent. That is, the objects which seem to be there in front of us are there in front of us. With hallucination, however, the relationship between ourselves and the objects our awareness is greatly obscured. It no longer seems adequate to explain our awareness as the straightforward consequence of an object's being there and our seeing it. In a sense, this means that there isn't a truly natural view of hallucination, and it places greater importance on the simple metaphysical picture to explain our awareness from that external, third-person perspective.

What is demanded here is an explanation of phenomenology in terms of some properties of the physical process that is taken to constitute the experience. Accordingly, we might explain phenomenology in terms of the *ways* in which certain constituents of the physical process affect the subject. If something is to feature in the phenomenology of experience we would expect that it must affect the subject in such a way that she should (assuming she is physically and mentally competent) be able to react appropriately and non-coincidentally to its perceptible features. Such appropriate reactions might include physical manipulations of the object, verbal descriptions of its perceptible qualities, artistic depictions of it, and so on. Since these behaviours are causally consequent upon physical states of the subject that are themselves caused by the object, it will be a requirement that the relevant states of the subject preserve and can transmit information about whatever perceptible features of the object are at issue.

Applied to the examples above, this requirement shows why, for example, the apple hallucination is of an apple and not (also) of an electrode: the electrode, for one thing, does not affect the subject visually, and the way in which the subject is affected by it is not one that preserves any significant information about any of the electrode's perceptible qualities. In the case of the subject's *seeing* an apple, the phenomenology of the experience does not feature the causally antecedent tree or seed because the apple's effect on the subject does not preserve any information about the visible qualities of the tree or the seed, only of the apple itself. It is therefore the nature and arrangement of only those qualities that can go on to causally influence the subject's apple-eating or apple-sketching behaviours. Note, however, that certain other constituents of the perceptual process *do* play a role in determining the experience's phenomenology: for example, the sun shining on the apple and the reflected light, although we might not say they are *seen* – they are not what we are looking at – nonetheless contribute to the phenomenology in terms of the brightness of the apple, its apparent colour and the various shadows and specular highlights that appear on its surface.

We might also note here that, in looking at the apple, we are (counterfactually) sensitive to some qualities of the sun – specifically its brightness – but not to others, such as its shape; it would make no difference to the phenomenology of our experience if the sun were cuboid instead of spherical. But it would make a difference if the sun were more or less bright. We are therefore sensitive to some features of the sun but perhaps not those (like shape) that incline us to say we *see* the sun. If nothing else, this suggests, perhaps helpfully, that there isn't a clear boundary between those constituents of the perceptual process that are perceived and those that are not. A clearer, but perhaps yet imperfect, boundary might lie between those constituents or qualities to which we are sensitive – which make a difference to phenomenology – and those to which we are not.

4.10 Conclusion - the argument from hallucination

I have in this chapter described the explanatory challenge presented by hallucination to proponents of the natural view, and various possible ways this challenge might be met. Three main strategies are available: (1) to acknowledge the genuine ('robust') nature of hallucinatory phenomenology and give a positive explanation of it in terms of awareness of something other than worldly objects; (2) to deny that hallucinations have robust phenomenology, and explain their *seeming* to have phenomenology in terms of their negative epistemic properties or cognitive effects; and (3) to acknowledge the robust phenomenology of hallucinations and explain it in terms of awareness of worldly objects. My aim in this chapter has been to show that the third strategy, although absurd at first sight, is tractable and perhaps even plausible given an account based on the simple metaphysical picture.

The simple metaphysical picture-based account has several crucial advantages over the other strategies. Indeed, it avoids most or all of the key problems that bedevil one or other of the alternative explanatory options. First, it acknowledges and accounts for the felt reality of hallucinatory experience in a way that equates felt reality with the possession of robust phenomenology. Second, it does so using an explanation whose essential ingredients are shared by the simple metaphysical picture of perceptual experience, making it explanatorily elegant and parsimonious. Third, by eschewing any appeal to awareness of something other than worldly objects it avoids the threat of the argument from hallucination while nonetheless sharing its pursuit of unifying explanation. Fourth, it has no difficulty accommodating cases in which the subject has insight into the hallucinatory nature of her experience. Fifth, in explaining the robust phenomenology of hallucination it allows us to explain the cognitive and behavioural effects of hallucination as being grounded in that phenomenology. And, sixth, by treating these cognitive effects as inessential accompaniments of hallucination rather than constitutive of it, the view allows that we can grant hallucinations to non-human animals and other subjects incapable of reflecting cognitively on their experiences.

The explanation of hallucination via the simple metaphysical picture also offered an explanation of several other sensory phenomena that might otherwise appear to threaten the natural view. As we shall see in the next chapter, it also points the way to an explanation of a widespread feature of ordinary perceptual experience, namely the way in

which its phenomenology is influenced by the subject's beliefs, judgments and expectations. This explanation will allow us to show that the natural view, while motivated by the seeming immediacy of perceptual experience, can also accommodate its *generality*.

CHAPTER 5 - PHENOMENOLOGICAL PROBLEMS (III): COGNITIVE INFLUENCES ON AWARENESS

5.1 Introduction

In the preceding two chapters I have described the challenges posed to the natural view by the phenomena of perceptual illusion and hallucination, and proposed some ways to respond to those challenges, in part using the simple metaphysical picture of experience. In this chapter I will discuss a further challenge arising from a perceptual phenomenon that features in many of our ordinary experiences, including the veridical ones, although it may also contribute to many cases of what we would call perceptual illusion. The phenomenon in question is the seeming role of cognitive states such as beliefs and desires in influencing the phenomenology of experience. Like hallucination, we seem to have here cases in which the phenomenology of experience is determined in part by factors other than the objects and qualities present in the scene before the subject. However, here it looks as though the extra factor is not some awareness of previously perceived objects but some cognitive processes.

In this chapter I will describe some evidence for the constitutive role of cognition in experiences' phenomenology, explain in more detail the problem it presents to a proponent of the natural view, and offer an account – again using the simple metaphysical picture – that positively explains the phenomenon while leaving the natural view intact.

Before that, however, it will be worth saying something more about the relationship between a subject's awareness of worldly particulars and her cognitive faculties. Although the motivating insight behind the natural view is the seeming *immediacy* of perception, we should be wary of neglecting a distinct but complementary aspect of perceptual awareness, namely its *generality*. That is, our perceptual encounters with worldly objects seem to present us with those objects in more than their bare particularity. Rather, we typically perceive those objects as objects of a certain *kind*. There is disagreement about the extent to which this generality is manifest in the phenomenology of experience, rather than issuing merely from some distinct judgment based on that phenomenology.⁵⁰

⁵⁰ For a defence of the constitutive role of kind-properties in phenomenology, cast in terms of representational content, see Siegel (2006).

We might usefully recall here Brewer's defence of naïve realism from the argument from illusion. He sought to protect the immediacy of perception by distinguishing the 'sensory core' of an experience from a cognitive component that grounds judgments based on that sensory core. The idea is that, in all cases, perceptual experience has this sensory core – constituted by the object itself and its perceptible qualities – but that in many cases, if not all, there is also the cognitive component which is responsible for acts of recognition and judgment. The upshot is what I have called the 'two-faculty' view that perceptual experience commonly involves the interplay of two more or less distinct faculties:

- 1. A faculty of (direct) awareness which acquaints the subject with the worldly objects and their qualities.
- A cognitive faculty that delivers judgments based on the objects and qualities revealed by the faculty of awareness.

As described earlier, the two-faculty view enables the naïve realist to hold that the sensory core of illusory experience remains attributable to the mere awareness of worldly particulars and their qualities, while the illusory aspect results from a more or less distinct act of judgment.

One thing the two-faculty picture lacks, however, is a clear account of how the two faculties are connected, i.e. how having things presented to her in perceptual experience constrains a subject's judgments about what state of affairs obtains. This question is a recurrent theme in the work of philosophers such as Sellars, McDowell and Davidson, where it is framed in terms of the capacity of our rational faculty – governed by holistic, inferential rules – to get a purchase on the causally-governed world of objects and events (see e.g. Sellars 2000: §36; McDowell 2009: 185; Davidson 2000: 156-7). My concern is chiefly to account for the manifest phenomenology of experience and not its epistemic status, so I will not seek to explore this question in great detail. Nonetheless, the relation between our sensory and conceptual faculties is clearly key to explaining the phenomenological role of cognition, which in turn presents a challenge to the natural view, as I will shortly describe.

Insofar as they do seek to account for the interface between perception and belief, naïve realists typically explain it in terms of a subject's capacity to recognise the things she perceives and categorise them as belonging to certain kinds of objects. Travis describes this as a subject's 'expertise' in judging what state of affairs obtains based on things' simply being that way (Travis 2007: 233-5). In other words, it is through the operations of what Fish calls a 'conceptual-recognitional capacity' that we are able to bring to bear on our beliefs the worldly things that we are simply presented with in experience (Fish 2009: 68).

Fish expands somewhat on Brewer's and Travis's account of this conceptualrecognitional capacity. His immediate concern is to account for an intuition that prior experience and knowledge can influence our perception of the world. He cites an example from Crane, who asks us to imagine the different experiences of a child and a scientist when presented visually with a cathode ray tube. The example is supposed to highlight both a commonality and a difference between the two subjects' experiences: a commonality in that both subjects see the same object; a difference in that only the scientist sees the object as being a cathode ray tube (Fish 2009: 68).

Fish explains the difference by recasting what it is that we see – not objects and/or properties, but *facts* understood as complex wholes of objects-bearing-properties; what are sometimes called states of affairs or situations (2009: 53).⁵¹ What the scientist possesses but the child lacks is the capacity to recognise the fact that is the object's being a cathode ray tube. And the scientist possesses this capacity in virtue of possessing the concept of a cathode ray tube; this is what makes it a *conceptual-recognitional capacity* (2009: 68).

As we saw earlier, Fish employs the same account to explain what he calls 'cognitive illusions', in which a subject perceives one kind of object but mistakes it for another (2009: 165 ff.). He uses the example of seeing a coiled rope and mistaking it for a snake. As he notes, such cases clearly cannot involve the subject's acquaintance with the fact of the object's being a snake, for the object is not a snake. Rather, the subject deploys her conceptual-recognitional capacity (for picking out snakes) inappropriately (2009: 169).

⁵¹ Fish is at pains to emphasize that this claim about what we perceive is intended as an expression of how the world seems to be presented to us rather than as a claim for the ontological priority of facts over objects and properties (2009: 53).

Fish acknowledges that this explanation only takes us so far before we are confronted with the disputed question of what constitutes a concept and hence what it takes to possess one (2009: 69). Without delving too deeply into this question, he suggests that an answer will "appeal in part to the highly specific interactions with the environment that the subject has enjoyed" (2009: 70).

It is important also to note that Fish's account of cognitive-recognitional capacities reflects something of his account of hallucination. Recall from chapter 4.5 above that Fish denies hallucination genuine phenomenal character and explains the 'felt reality' of hallucination as mistaken beliefs brought about by the (non-phenomenal) hallucinatory state (2009: 94). This solution allows Fish to accept that hallucinatory states are wholly internal (they supervene locally on neural states so that these neural states are sufficient for the occurrence of those hallucinations) while avoiding the pressure to say the same about veridical perceptual experiences. It seems that he also takes the cognitive factors involved in conceptual-recognitional capacities to be internal to the subject (2009: 173). Putting the two together, we can conclude that the employment of a conceptual-recognitional capacity should not alter the phenomenology of a perceptual experience, but instead influence the beliefs and judgments we form on the basis of that experience. Although the clear distinction between external (perceptual) and internal (cognitive or conceptual) factors might appear somewhat blurred by Fish's above-noted claim that perception acquaints us first and foremost with facts, we should recall that the relevant facts are to be understood as the objective ways things are arranged perceptibly before the subject. This allows Fish to develop a disjunctive approach to veridical versus illusory cases of fact perception, so that the fact in the former case is a genuine contributor to the character of the experience, while the subject in the illusory case merely and mistakenly believes there to be a fact 'out there' and hence that there is the phenomenal quality that would be constituted by that fact (2009: 171).

The same worry therefore applies to Fish's account of cognitive illusion as to his account of hallucination, namely that it fails to do justice to our sense of the 'felt reality' of the illusory phenomenology. As with hallucination, however, Fish needs to avoid positing a genuine effect on phenomenology from such cognitive effects, on pain of undermining the unique claim of what is 'out there' to 'shaping the contours' of experience.

In this regard, we might note it is implicit in Brewer's description of "two levels in the subjective character of experience" that for him both the simple awareness of the object and the application of the relevant concept play a role in determining the phenomenology of the experience. This is confirmed explicitly in more recent work, where Brewer characterises the two levels in terms of the object's 'thinly' and 'thickly' looking some way to a subject: it *thinly* looks a certain way F in that it in fact looks relevantly similar to paradigm instances of F things; and it *thickly* looks F if, additionally, the subject recognises it as an instantiating F-ness (Brewer 2011: 121). Brewer insists that this application of the concept of F-ness – one's seeing the object thickly as an instance of Fness - makes a genuine difference to the phenomenology of the experience, but without altering how the object *thinly* looks (2011: 123). So it seems that the respective contributions to phenomenology of something's looking thinly and thickly F remain independent of one another at least insofar as how it *thickly* looks does not alter how it thinly looks. Presumably it is important for Brewer's naïve realism to maintain some distinction between, as we might put it, perceptual and conceptual phenomenology, for fear of leaving the view vulnerable to the generalising step and inviting the challenge to explain why all of the phenomenology of experience shouldn't be determined by internal states of the subject.

There is compelling evidence, however, for the claim that our sensory and conceptual or cognitive faculties are much more closely intertwined than the two-faculty picture suggests, so that our acts of recognition or belief do not merely take the phenomenology of experience as an input but can actually influence the phenomenology itself. On the face of it, this would seem to present a challenge to the natural view, since any admission that ostensibly 'inner' processes of recognition contribute to phenomenology once again challenges us to explain why we shouldn't think that *all* phenomenology is determined by processes internal to the subject. I will, however, go on to offer an explanation of this cognitive influence on phenomenology which is consistent with the natural view and in which the notion of a conceptual-recognitional faculty plays a central role. Unlike Fish's account, but in keeping with my earlier account of hallucination (chapter 4), my account will attribute the phenomenology, both perceptual and conceptual, to awareness of worldly objects.

5.2 Perceptual learning and cognitive penetration

There are two more or less distinct phenomena that offer evidence for the genuine effects of cognitive factors on the phenomenology of perceptual experience. These phenomena are *perceptual learning* and *cognitive penetration of perceptual experience*, and I will consider them in turn in this section. In the following section I will describe some evidence that one or both of these phenomena involve awareness of qualities not attributable to the scene perceived.

Perceptual learning is the process whereby repeated perceptual exposure to some object or type of object results in a more or less permanent change in one's perceptual experience of the same (Gibson 1963: 29). Perceptual learning is responsible for the ability of subjects to improve their discriminatory powers with respect to certain kinds of perceived objects and qualities, sufficiently in many cases for those subjects to become 'expert perceivers' such as wine tasters and field botanists. Whether perceptual learning poses a challenge to the natural view depends on what sort of mechanism we think is responsible. Gibson notes that theories of perceptual learning fall into two broad categories: what she calls *enrichment* theories, whereby our prior experience (including memories and beliefs) *adds* something to our current perceptual experience; and *differentiation* theories, whereby we become better able to attend to and discriminate certain of the more distinctive or salient qualities of perceived objects (Gibson 1963: 40; see also Gibson & Gibson 1955).

If differentiation theories are correct and perceptual learning elicits changes to the phenomenology of experience by modifying attention or focus in some way, then there is not obviously a problem for the natural view. A change in focus or attention does not introduce any phenomenology that cannot be attributed to the scene before the subject; it merely alters the way in which we single out, or indeed group, certain objects and qualities for particular attention. We can allow that this amounts to changing the phenomenal character of the scene, insofar as phenomenal character is understood merely as that relational property of a scene which is how it appears perceptually to a subject. Any change in phenomenal character here is perhaps best understood as a change in *how* we are aware of a scene; if there is an influence on *what* we are aware of insofar as we are more or less likely to notice certain qualities or groupings of qualities, there is nonetheless no

134

change in the qualities that are presented to us and therefore no change in what we *can* be aware of.

Fish's claim that we perceive facts is pertinent here (recalling that a fact for Fish's purposes is to be understood as an obtaining state of affairs rather than a true proposition) (Fish 2009: 52-3). Fish claims that we can allow for a difference in phenomenology between experiences of expert and inexpert perceivers, and that we can account for this compatibly with naïve realism by appealing to the different 'conceptual-recognitional capacities' of the different subjects (2009: 68ff. & 167ff.). Expert perceivers recognise that certain arrangements of qualities belong to a certain kind of object while inexpert perceivers might be aware of all the same qualities without recognising that the relevant fact obtains (for example, that the mottled patch in the leaf litter is a woodcock). We might seek an explanation for the development of recognitional capacities via a differentiation theory whereby prior perceptual encounters with a certain kind of object alter the salience of certain qualities or groupings of qualities and so influence the way in which subjects attend to or focus upon the object. Crucially, this would not require a change in what the subject is aware of - which objects and qualities are presented to her but a difference in the extent to which different elements of the scene 'stand out' in awareness.

The second kind of perceptual learning theory appears less easy to accommodate to the natural view. Enrichment theories claim that repeated exposure to a type of stimulus alters the subject's successive experiences not (or not merely) by altering her attentional focus and her ability to discriminate certain qualities or groupings of features but by adding something to the later experiences that alters or embellishes their phenomenology. We might, for example, think of the added ingredient as memory images which are evoked by the stimulus through a process of association (Gibson 1963: 40; Gibson & Gibson 1955: 34). As Gibson & Gibson put it, the result is that with increasing experience and expertise "perception is progressively in *decreasing correspondence with stimulation*" (1955: 34); or, in terms of the natural view, we might say that the phenomenology becomes less attributable to the object before the subject and more attributable to remembered objects. If enrichment theories are true then there may be more to perceptual phenomenology than is contributed by awareness of the worldly objects and qualities perceived, in which case

we are confronted again with the generalising worry that was threatened by the arguments from illusion and hallucination.

Differentiation and enrichment theories are, of course, not mutually exclusive – there is no reason why both mechanisms might not be operative in some cases, even if others involve only one. Ornithological expertise and familiarity with hen harriers might enable me to distinguish them from marsh harriers in part by focusing on salient features (a whitish rump present on the former but not the latter), but it may also augment the phenomenology of my experience by making the whiteness appear even whiter than can be accounted for by the qualities inhering in the bird itself. Deciding which mechanism is operative in any specific case may be notoriously difficult (Gibson 1963: 40ff.) but, crucially for our current concern, there is good evidence that something like enrichment can occur in at least some cases. I will consider this evidence in the next section.

While perceptual learning is considered to involve a more or less permanent change in the perception of some target class of stimuli through repeated exposure (Gibson 1963: 29), presumably by some associative process beyond the subject's conscious control, there is another way in which prior experience might influence the phenomenology of perception. Thus, there is good evidence that a subject's perceptual experience can be shaped by her concepts, beliefs or desires concerning the object perceived, all of which may be consciously brought to bear on the experience of that object. It is said that, in this regard, the subject's perceptual experience is *cognitively penetrable* by the subject's beliefs, etc. What the thesis of cognitive penetration entails can be understood by way of its opposite, namely the claim that perceptual experience is cognitively *im*penetrable:

[P]erceptual experience is cognitively impenetrable if it is not possible for two subjects (or one subject at different times) to have two different experiences on account of a difference in their cognitive systems which makes this difference intelligible when certain facts about the case are held fixed, namely, the nature of the proximal stimulus on the sensory organ, the state of the sensory organ, and the location of attentional focus of the subject.

Macpherson 2012: 29

Note that, to count as genuine cognitive penetration, the alteration to phenomenology should reflect the *content* of the influential cognitive state. This rules out indirect effects

of cognitive states on phenomenology, such as, say, visual disturbances brought on by fearful beliefs (see Macpherson 2012: 26).

5.3 Empirical evidence for cognitive contributions to phenomenology

Several studies have shown that subjects' memories and expectations exert an influence on their perception of colour. In a classic study by Delk and Fillenbaum (1965), subjects were exposed to a set of nine shapes cut from card of identical, uniform colour (orange-red). The nine shapes fell into three groups according to their conventional association with the colour red: the first group (heart, apple and lip shapes) are positively associated with redness; the second group (oval, circle and ellipse) are neither positively nor negatively associated with redness; and the third group (horse, bell and mushroom) negatively associated with redness. The shapes were presented against a background composed of a colour-mixer that could be adjusted by the subjects to any shade between red and yellow-orange. Subjects were asked to view each shape in turn and adjust the colour-mixer until it was the same colour as the shape, such that the shape became indiscernible from the background. The crucial finding was that, although all of the shapes tested were cut from the same uniformly orange-red card, subjects tended to adjust the background to different colours depending on which shape was being viewed. Where the three conventionally red shapes were presented, the subjects tended significantly to adjust the background to a redder shade in order to achieve the subjective match than was required for the three neutral shapes. A more limited effect in the opposite direction was achieved in respect of the shapes negatively associated with redness, i.e. in some cases there was a significant tendency to select a more yellow-orange background than for the neutral figures (Delk & Fillenbaum 1965: 292-3).

The authors conclude that previous experience of shape/colour associations influences perceived colour. Hansen et al. (2006) draw a similar conclusion from their own experiments, in which subjects were shown digital photographic images of familiar fruits presented against a uniformly grey background and asked to adjust the colour of the fruit images until they appeared grey, i.e. lacking chromaticity (colour) but still with normal variation in luminance (brightness). The researchers found that subjects tended to overcompensate in their colour adjustment in order to see the fruits as grey, e.g. they adjusted the banana image until it was objectively slightly blue (ibid.: 1367).

Concomitantly, when the banana image was objectively grey (achromatic) the subjects reported it as being slightly yellow (ibid.). The authors conclude that "knowledge of the world affects our perception" (ibid.: 1368).

It is not obvious whether these two cases ought to be considered evidence for perceptual learning, on the one hand, or cognitive penetration on the other. Macpherson presents the Delk and Fillenbaum experiment as likely evidence for cognitive penetration, but acknowledges that the results might instead be explained in terms of the influence of past experience (of characteristically red shapes) via some associative mechanism, and without the input of concepts, beliefs and so on (Macpherson 2012: 45). In other words, these might be cases of perceptual learning rather than cognitive penetration. As Macpherson notes, this would help to explain the likely imperviousness of the effect to changes in the subjects' beliefs (ibid.: 46).

Macpherson offers as a more clear-cut case of cognitive penetration an experiment by Levin and Banaji (2006), in which subjects were shown various greyscale images of faces which were computer-generated to have stereotypical features of black people in some cases and stereotypical features of white people in others. Crucially, the images could also be adjusted in respect of their luminance and contrast, so that image pairs of 'white' and 'black' faces could be produced in which both images were of the same mean luminance and contrast. In one experiment, subjects were presented with such an image pair, and asked to adjust the luminance of each in turn until it seemed to match the luminance of a further reference image, which was again of a 'black' or 'white' face, but possibly of some different starting mean luminance. Where subjects had to match an adjustable 'black' face with a fixed reference 'white' face, they would tend to over-adjust the luminance of the adjustable image, i.e. make it even lighter than the reference image. Conversely, they would tend to adjust a stereotypically 'white' face to make it too dark when matching against a 'black' reference image (Levin & Banaji 2006: 504).

This suggests that the appearance of stereotypically 'white' or 'black' facial features makes the images look, respectively, lighter or darker than they really are. As described, the experiment is inadequate to settle whether this is due to cognitive penetration or perceptual learning, since it could be that the perceived link between lightness and other facial attributes is merely associative, even if this seems a more sophisticated association than between, say, heart-shapes and redness.

That cognitive penetration is involved seems clearer, however, from a further experiment performed by the same authors. In this test they again presented subjects with image pairs, but this time one of the images (for half of the subjects the 'white' face and the other half the 'black' face) was replaced with an 'ambiguous' face whose features were, roughly speaking, intermediate in form between the stereotypical 'black' and 'white' faces. The members of each image pair were labelled as 'white' or 'black' so that each 'ambiguous' face was labelled unambiguously as 'white' when its pair was 'black' and *vice versa*. This time, the subjects were asked to match each member of the image pair in turn to an adjustable grey rectangle. The result was that subjects adjusted the grey rectangle to a darker shade when matching against an 'ambiguous' face labelled as 'white' (Levin & Banaji 2006: 505-6).

Unlike the cases presented by Delk and Fillenbaum and by Hansen et al., this experiment clearly shows that the nature of the perceptual effect is to some extent under the control of the subjects' conscious beliefs. Unfortunately, it is at the same time less clear that Levin and Banaji's results couldn't be explained merely as the result of differing judgments by the subjects involved. What convinces that there are real phenomenological differences in the other studies – that the heart-shape really does *look* redder, and that the grey banana image really looks yellowish – is the way in which the images are presented against, and indeed embedded in, their background. By leaving the target and reference images side-by-side but separate, Levin and Banaji leave room to attribute any mismatch to judgment rather than a real difference in the phenomenal qualities of which the subjects are aware.

This shortcoming might be remedied by an amended methodology, but what matters for my account is that at least one of perceptual learning or cognitive penetration demonstrably involves some aspects of phenomenology that cannot be attributed to the scene before the subject. This poses a similar challenge to the natural view to that posed by the kinds of illusory and hallucinatory phenomena described in previous chapters. In the following sections I will argue that the simple metaphysical picture allows us to explain this cognitive influence on phenomenology along the same lines as the explanation of hallucination offered in chapter 4. The mechanism is perhaps most easily expressed in respect of perceptual learning, but I will go on to show how it might be extended to accommodate cognitive penetration.

5.4 A simple metaphysical picture of perceptual learning

I have suggested that in at least some cases of illusion, the argument can be avoided by insisting that, in fact, the phenomenology of the experience *can* be exhaustively attributed to the scene itself if we allow for the constitutive role of factors such as ambient lighting. Hallucination I admitted cannot be accounted for in this way, but I have suggested instead an alternative explanation (consistent with the natural view) involving the constitutive role of previously seen objects in the phenomenology of current experience. Now, with perceptual learning and/or cognitive penetration, we appear to have cases like hallucination where phenomenology cannot be attributed exhaustively to the scene arrayed before the subject. In this case, however, the 'extra' elements of phenomenology – additional to what is constituted by the scene before the subject – seem attributable to a subject's prior experience, expectations or beliefs.

This claim about the constitutive role of cognitive factors in the phenomenology of experience can be used to form a variant of the argument from illusion/hallucination. Thus, if some part of an experience's phenomenology can be determined or constituted by cognitive factors, and if we further think of these cognitive factors as in some sense internal to the subject, then why should we not allow that the *whole* of that experience's phenomenology is determined or constituted by more or less internal states of the subject?

In this section I will propose a reply to this argument that uses the natural view to show how, even in cognitively-influenced experiences, we might attribute the phenomenology wholly to awareness of worldly things. The proposed mechanism has the following virtues:

 it acknowledges the inseparable cognitive or conceptual aspect of much or all perceptual experience;

- it offers a unified explanation of both the sensory and conceptual elements of perceptual experience and so resolves any puzzle about how the two combine to yield an experience's overall phenomenology;
- it holds that the conceptual elements, like the sensory elements, are 'out there' and so avoids the potentially problematic implications of phenomenal qualities arising from purely inner mental states or factors.

My suggested account will draw on my earlier account of hallucinatory experience, and specifically the claim that the seemingly problematic elements of phenomenology derive from previously perceived objects (see chapter 4). The aim of that account was to explain how hallucinations come about in a way that does justice to the intuition that they possess 'robust' phenomenology. The upshot was that we can account for hallucination consistently with the natural view, as having genuine phenomenology *for the same reason that perceptual experiences have phenomenology*, namely because hallucinations too are object-involving experiences and so do not supervene on neural states.

I arrived at this surprising conclusion via a parallel between a certain kind of hallucination and certain (allegedly problematic) cases of perception. The perceptual cases in question were those involving substantial 'time gaps', i.e. a marked lapse of time between the event perceived and the act of perceiving it. Such time-gap cases merely reveal what, given the natural view, is always true to a greater or lesser extent, namely that past events are constituents of current experiences.⁵² I suggested that hallucinations could be understood as the reappearance or re-presentation of objects previously perceived. In light of the simple metaphysical picture, this reappearance was understood as the recurrence in the subject of a similar neural state as that caused by the object when it was originally perceived, in such a way that the current neural state is *also* relevantly an effect of the same object (i.e. the same temporal stage of object as previously perceived). Of course, the object exerts its effect on the subject's *current* neural state by a more convoluted and temporally drawn-out route than that by which it exerted its original effect in being perceived. Nonetheless, given the simple metaphysical picture and the time-gap

⁵² Cases involving extreme time gaps, e.g. a subject's perceiving distant celestial events that took place thousands of years before her birth, are typically cited in order to undermine claims for the directness of perception (e.g. Robinson 1994: 80-4). I sought to avoid the *reductio ad absurdum* by noting that it seems compelling only because of an antecedent commitment to a certain substantial notion of 'experiences' as states or events that in some sense 'belong' to the subject.

cases of perception, there seems no principled way to rule out such convoluted and drawnout processes as cases in which the object contributes constitutively to an experience's phenomenology.

How, then, does this account of hallucination bear on the proposed mechanism for perceptual learning? The next ingredient of the account comes by way of two potential objections to the natural view and simple metaphysical picture.

The first objection is one I alluded to in the earlier discussion of hallucination, namely that the process involving the perceived object and the perceiving subject is open-ended historically and has an infinite or near-infinite number of constituents besides the object that ostensibly is presented in experience. The challenge to the proponent of the natural view and simple metaphysical picture is to explain just why it is that object and not any other constituent that contributes constitutively to the phenomenology of the experience. My answer was that what distinguishes the perceived object from (some) other constituents of the perceptual process is that the relevant neural state of the subject preserves a substantial amount of information about the perceptible qualities of that object. Furthermore, we should expect that this information-preserving causal relation is one that persists but is updated over time – that changes in the subject's neural state are sensitive to changes in the perceptible qualities of the object. This continuing sensitivity that can be understood as *tracking* (whether we think of this as the neural state's tracking the object or its making possible *our* tracking of the object).⁵³

The second potential objection is that a given perceptual neural state must inevitably be the product of a variety of causes that are themselves not directly causally related. This multiplicity of causes is especially marked given some further assumptions:

⁵³ We might worry here that the subject's neural state effectively tracks all of the intermediate parts of the perceptual process between the object and the subject, since each intermediate part must equally carry the information that is eventually transmitted to the subject. However, what distinguishes the perceived object from later time-slices of the perceptual process is that the former is the *origin* of the information or what, in a sense, the information is *about*. The subject's neural state in a sense tracks all time-slices between the object and itself, but the object is what it *ultimately* tracks; all other intervening time-slices of the process are tracked *in virtue* of the neural state's tracking the object. Meanwhile, the further back we go prior to the relevant temporal stage of the object, the less information the neural state contains about those earlier time-slices.

- 1. There is considerable overlap in the neural locus of perceptual effects brought about by (possibly greatly) different perceived objects.
- Some of the neural effects at these loci are persistent beyond the duration of the relevant perceptual episode, and these persistent effects ('traces', we might call them) are at least partially responsible for subsequent conscious recollection of the object as originally perceived.
- 3. The traces left by an object A at time t may constrain the perceptual effects wrought by a different object B at a later time (t + 1) insofar as the neural loci of A's traces and B's effects overlap. Where this is so, the neural effect of B is also to some extent an effect of A.

The overlap suggested in (1) may be reduced by a degree of neural localisation, but the scope for this must be limited given, on the one hand, the finite capacity of the relevant parts of the brain and, on the other, the boundless diversity of possible objects of perception.

Given the third assumption, adherence to the simple metaphysical picture might encourage the thought that, in seeing B, the subject is also in some sense 'seeing' A. Apart from anything else, the partial nature of the overlap between the neural loci of A and B's perceptual effects (assuming that A and B are perceptibly different) means that only some parts or elements of B's overall effect are constrained by the persistent effects of A. It is easier to segregate the effects of B (which we see) from those of A (which we have previously seen) if we think in terms of perceptual *tracking*. Arguably, seeing an object necessarily involves an ability visually to track changes in that object in real time, so that this tracking relationship is partly constitutive of what it is to *see* an object. A crucial distinction between hallucinated and seen objects, then, is that we are tracking the latter but not the former. Returning to the simple metaphysical picture of cognitive penetration, we might talk more loosely of a subject's *neural states* ' tracking an object, where this denotes a sensitivity or causal correlation between changes in those neural states and changes in the object or our perspectival relation to it.⁵⁴

⁵⁴ When we cast it in neural terms, it is easier to see that our visually tracking an object is not *sufficient* for seeing it; consider seeing a television, where the activity of one's visual system is causally correlated with some internal workings of the television, internal workings that are certainly not seen. In addition, more
With these assumptions and caveats in mind, we can now turn to the proposed mechanism for perceptual learning. As a start, we might endorse the naïve realists' appeal to a conceptual-recognitional capacity as the 'entry-level' interface between perceived objects and our thoughts and beliefs. At the most basic level, what this capacity surely consists in is sensitivity to perceptible similarities amongst distinct particulars, which similarities ground categorisation. The model I am proposing asserts that the engagement of a conceptual-recognitional capacity in perception consists, at the most rudimentary level, in the perceived object's evoking or bringing to mind other previously perceived objects that are relevantly similar and which are therefore likely to be assigned to a common category or kind. What I mean here by 'evoking' and 'bringing to mind' will hopefully become clear, but it involves those previously perceived objects becoming available to awareness.

Before attempting to defend this surprising conclusion, I should say a little more about the role of perceptible similarity in the account. I take it that a conceptual-recognitional capacity exploits perceptible similarities between objects of the same type or temporal stages of the same object. Furthermore, I assume that the exploitation of similarities among objects or temporal stages of an object is made possible in turn by similarities among the *perceptual effects* of those objects or temporal stages, notably their effects on the subject's neural states.

This latter appeal to similar perceptual effects avoids a concern about circumscribing the range of possible similarities among perceived objects. The concern is that one could posit all manner of similarity relations between any pair of objects; claiming that we are perceptually sensitive to objective similarities therefore invites the demand that we explain just which of many possible similarity relations we are sensitive to. Appealing instead to similarity relations among perceptual effects, rather than what these are effects *of*, obviates the need for any such explanation.

The requirement for similar perceptual effects is also made plausible by the thought that a subject's application of the same concept to perceptual experiences of different

needs to be said to generate a *necessary* condition for seeing in terms of tracking. For example, we can see things for a split second but it is questionable whether we can be said to track things over such a short period; the condition should perhaps be that we can track them for as long as they affect us visually.

objects or temporal stages of an object is itself a common kind of effect of diverse experiences. It might be objected, however, that the same problem recurs in respect of similarity among neural effects. Which of many possible similarity relations among neural states are significant in relation to the respective perceptual experiences or the application of a concept to both?

Implicit in the question is the thought that we must, in applying a concept, select from among various possible similarities among, on the one hand, the perceived objects or, on the other, their neural causes. The solution, however, is to recognise that what is involved in bringing diverse objects under the same concept is a certain commonality of behavioural response to those objects. Just what those behavioural similarities are will depend on the concept under consideration, but the principle can be illustrated by example. So, applying the concept CAT to diverse perceptual experiences of cats (and misleadingly cat-like non-cats) will in each case consist in, for example, the subject's assenting to the utterance 'Is that a cat?', reaching for the cat food, and so on.⁵⁵

To see how our conceptual-recognitional capacity involves the exploitation of similarities among diverse objects, and how my account uses this to posit a constitutive role in *current* perceptual experience for *previously* perceived objects, it will help to consider an example. Let us imagine a subject's sequence of perceptions of a series of similar objects – the subject's seeing some sunflowers, say. According to my earlier assumption, the similarity of the environmental stimuli will be reflected in certain similarities in the resulting perceptual neural states; not least a similarity in the locus of neural effect.⁵⁶ Again, let us assume that the subject's visual experience leaves 'traces' at the locus of its perceptual effect, traces which underwrite later recollection of the object and its perceptible qualities. It follows from these two assumptions that the subject's seeing a sunflower at time t1 (let us call this object sunflower 1) will leave a trace over more or less the same neural locus that is activated by a sunflower seen at a later time t2

⁵⁵ Whether we could enumerate necessary and/or sufficient conditions for possession of the concept CAT is debatable and arguably beside the point. As will become apparent, an account of concepts relevant to the sorts of perceptual phenomena considered here will take the notion of concept possession as having to do ultimately with subjects' propensity to assimilate certain worldly objects and qualities in virtue of how they appear perceptually.

⁵⁶ I am ignoring possible differences in the subject's perspective on the different sunflowers and possible differences in background conditions such as illumination. These possible differences potentially complicate a theory of concept formation along the lines I am proposing, but it is beyond the scope of the current thesis to address this.

(sunflower 2). Now, given my assumption (3) above, this trace will to some extent influence or constrain the neural effect wrought by sunflower 2. In other words, this neural effect would have been different had the subject not previously seen sunflower 1. What this means is that the neural state instantiated in seeing sunflower 2 is the effect of both sunflower 2 and sunflower 1, and furthermore that the neural state in some measure carries information about, because it is isomorphic to, both sunflowers 1 and 2. As the subject goes on to see sunflowers 3, 4, 5 and so on, there develops a layering or progressive modification of traces over roughly the same neural locus, so that each subsequent neural state is causally constrained by all of those that preceded it.

Now, when we bring the simple metaphysical picture back into the frame, it might seem that sunflowers 1, 2, 3, 4, etc. are all on a par when it comes to seeing the latest sunflower in the series. All of them participate causally in the perceptual process culminating in the subject's instantiating the relevant neural state. And in each case the form of the sunflower influences the form of the neural state such that there is a significant degree of isomorphism between one and the other. We might therefore wonder whether it follows, given the simple metaphysical picture, that the subject's seeing sunflower 4 is also, to some extent, an experience of sunflowers 1, 2 and 3.

On the face of it, this is implausible. The kind of experience we wish to account for is explicitly an ordinary one in which it appears to the subject that she perceives only the one sunflower – the one that is before her at the time, namely sunflower 4. Nonetheless, the conclusion I wish to draw is precisely that sunflowers 1, 2 and 3 play a constitutive role in the experience of sunflower 4 in such a way as to influence its phenomenology. Clearly, however, the four objects are *not* on a par when it comes to the experience and its phenomenology. What is needed is an explanation of how the other sunflowers exert an influence over the experience's phenomenology and how this differs from the influence of the one that is seen.

All of sunflowers 1-4 affect the subject causally and appropriately via her visual system, albeit that the causal process is, for sunflowers 1-3, mediated by some 'memory traces'. Furthermore, we might assume that the causal process is in each case appropriately information preserving, i.e. the effect (the subject's perceptual brain state) is one which contains significant information about the visible properties of each of

sunflowers 1-4. As I suggested earlier in this section, the crucial difference between the subject's awareness of sunflower 4 and her awareness of sunflowers 1-3 is not that the former lacks a mediating memory trace, but that this awareness permits her to track sunflower 4 over time through changes in its perceptible properties. To put it in terms of the physical process of perception, the subject's perceptual neural state comes under the causal constraint of sunflower 4 in a way that it doesn't to sunflowers 1-3, i.e. there is a causally-mediated correlation between changes in sunflower 4 and changes in the subject's relevant neural activity.

What, then, should we say of sunflowers 1, 2 and 3 if not that they are seen? Nor should we say that they are *recalled*, despite the fact that their role in a current experience is akin to that of a remembered (or hallucinated) object, inasmuch as they are previously perceived objects playing a constitutive role in a current experience, and so coming back into awareness. After all, it is not typically the case that, in seeing sunflower 4, a subject will *consciously* recall any or all of sunflowers 1, 2 and 3. In other words, sunflowers 1, 2 and 3 do not *obviously* intrude into the phenomenology of the experience.

Nonetheless, I have claimed that sunflowers 1, 2 and 3 *do* in fact contribute constitutively to the phenomenology of the experience of seeing sunflower 4. So how is this manifest if it is typically not obvious to the subject upon reflection? Recall the experiments that revealed the cognitive penetration of perceptual experience. Presumably in these cases – of heart-shapes appearing redder than their objective hue, bananas more yellow, etc. – the subjects were not reflectively aware of these biases in their perceptual experiences until they were revealed to them by the experimenters. Nonetheless, the effects were demonstrably real.

In the case of the heart-shapes, my theory would suggest that the 'extra' redness not objectively qualifying the *currently seen* heart-shape is constituted by the actual redness of *previously seen* heart-shapes. A notable feature of this phenomenon is that these previously seen heart-shapes and their qualities are manifested in the experience's phenomenology 'merged' with the currently seen shape and its qualities. In other words, we do not see the heart-shape located before us and have some vague and unlocated sense of 'extra' redness – this redness qualifies the same apparent location as the genuinely seen shape. The perceived object is thus, in a sense, imbued or amalgamated with more or less

faint 'traces' of similar objects previously perceived. It is presumably a combination of this 'faintness' of the previously perceived objects and their subsumption within the apparent location of the currently perceived object that accounts for our usual lack of reflective awareness of cognitive penetration.

There are thus two aspects to what we might call the 'domination' of the currently perceived object over the previously perceived objects: first, the former's determination of the perceived location of the relevant phenomenal qualities; and, second, the greater contribution of the former's qualities to the phenomenology of the experience. This second factor might be explained in terms of the magnitude of their respective causal influence on the perceptual neural state. We can understand this in counterfactual terms: if we take a set of objects that contribute causally and formally to a neural state N, and consider counterfactual situations in which each object in turn is deleted from the relevant causal history, the counterfactual neural states N', N'', N''' etc. will each differ to a greater or lesser extent from the actual state N. The limit case will be the counterfactual situation in which the currently perceived object is deleted (Sunflower 4 in the previous example): in this case neither neural state N nor anything relevantly similar will be activated. Here we can see a special sense in which the currently perceived object plays the dominant causal role in bringing about N. There is an asymmetric causal dependence of the effects of sunflowers 1 to 3 on the effects of sunflower 4: sunflowers 1 to 3 would not have effected any current neural activity at the relevant locus were it not for the effects of sunflower 4, whereas sunflower 4 would still have wrought some such (if somewhat different) effects had some or all of sunflowers 1 to 3 never been seen.⁵⁷

I have described the effect of perceptual learning as one in which the qualities of previously perceived objects are in a sense amalgamated with the qualities of whatever is currently before the subject. This implies a sort of *superimposition* of the past quality on the present. Note, however, that the case of the banana image (Hansen et al. 2006) shows that we would be wrong to think that the cognitive effect literally involves the superimposition of the colour of past bananas (yellow) over that of the current banana-image (blue-grey), for that would presumably result in the appearance of greenish grey rather than grey. Rather, the effect appears to be one of *interference* more than

⁵⁷ This notion of asymmetric dependence owes something to Fodor's account of perceptual content (Fodor 1989: 106-110).

superimposition, at least in this case where, unlike the Delk and Fillenbaum example, the contribution of the cognitive faculty is a colour complementary (i.e. opposed) to that possessed by the target object.

5.5 From perceptual learning to cognitive penetration

In summary, then, the proposed theory of perceptual learning, drawing on the simple metaphysical picture, holds that the phenomenology of experience is attributable both to awareness of what is presently perceived and to awareness of one or more things previously perceived. Those previously perceived objects that contribute significantly to phenomenology in this way are 'brought to mind' by virtue of their perceptually relevant similarities to the object currently perceived.

As far as it goes, this might sound like a possible account of perceptual learning, since the influence of previously perceived objects is more or less permanent (as permanent as the memory 'traces' presumably) and not obviously susceptible to conscious control by the subject. It would be reasonable to consider the mechanism genuinely *cognitive* inasmuch as it purports to form the basis of conceptual-recognitional capacities. All the same, there is no explicit role for beliefs or similarly sophisticated cognitive states in generating the influence, so this is clearly not suitable as an account of cognitive penetration.

Extending the account to accommodate cognitive penetration might, however, be achievable without too much modification. Consider the Levin and Banaji experiment with the ambiguous face images (and assume that this is genuinely a case of cognitive penetration). Here the subjects' beliefs about the ambiguous image are engineered by labelling the images as either 'black' or 'white'. Perhaps what is going on here is that, in the given context, the written word 'black' evokes a 'memory image' of one or more black persons' faces. In line with my account of perceptual learning, the 'memory image' is really constituted by awareness of those previously perceived faces. As in the heart-shape or the banana-image, the recalled objects (faces in this case) are 'superimposed' on the seen image such that it seems to the subjects that they are aware of only the one object. The difference with perceptual learning is just that the recollection or evocation of previously perceived objects is augmented or potentiated by some stimulus other than the target object.

It is likely, of course, that the ambiguous face image itself evokes previously perceived faces, albeit that the ambiguity is likely to evoke previously-seen black and white faces more or less equally. What the addition of the label does is to skew the balance in favour of recollecting black faces. It might be assumed that the provision of the label here serves to prompt a conscious belief of the sort expressed by 'That is a black face', although it is unclear, given the proposed mechanism, whether we really need to think of this belief as a genuine intermediary between the image and the phenomenological effect. Certainly, we need some way of accounting for the way in which the unqualified word 'black' is understood to refer to the character of the represented face. Furthermore, we need a mechanism by which perception of the written word 'black' inspires recollection of previously perceived black faces. Some sort of associative mechanism seems likely here, whereby the brain effects appropriate causal linkages between perceived words (e.g. the written or spoken word 'black') and certain objects. A given word might become associated with different objects for different reasons, so that ambiguity results (think of black people, black cats, black moods, etc.), and a plausible associative theory would need to account for the role of context in effecting disambiguation, presumably by providing other perceptible objects that reinforce some associative linkages and not others. A detailed account of this sort is beyond the scope of this thesis, but we might begin to see how the notion of beliefs as intermediary states or events between an object perceived and one (or more) recollected might not be necessary. Just as the associative mechanism proposed to account for perceptual learning might form the basis for what we call a conceptual-recognitional capacity, the evocation of previously perceived objects (what we might metaphorically call the 'priming of memory images') by certain external cues (written words, for example) might be sufficient in itself to account for what we call the influence of belief on perception in at least some cases.⁵⁸

⁵⁸ Compare Macpherson's 'indirect mechanism' of cognitive penetration, which takes the latter to involves two steps: first, the relevant cognitive state by a quasi-imaginative process generates some appropriate phenomenal qualities (as when, for example, the belief that an image is a heart shape evokes the colour red and puts it before the mind's eye, as it were); and, second, this phenomenal quality is in some more or less literal sense amalgamated with the phenomenal quality instantiated in virtue of one's perceptual experience of the heart-shaped image (Macpherson 2012: 54-5).

5.6 Ambiguous figures and error

Cognitive penetration has been held responsible for the familiar phenomenon of 'aspect shifts' involving ambiguous figures, as when the famous duck-rabbit figure is seen at one moment as a duck and at another moment a rabbit. This phenomenon might be accommodated using the same mechanism as described in the previous two sections. Assuming that the figure's ambiguity comprises its roughly equal similarity to both stereotypical ducks (or duck images) and stereotypical rabbits (or rabbit images), and that this similarity is reflected in a similar degree of congruence or overlap between the following pairs of neural states: [Nf (neural state caused by the figure) & Nd (neural state caused by typical duck/duck image)] and [Nf & Nr (neural state caused by typical rabbit/rabbit image)]. With such a similar degree of congruence between Nf and Nd as between Nf and Nr, we might suppose that there is a more or less equal likelihood of the figure evoking previously seen ducks as rabbits (in other words, a similar likelihood that the subject will see the figure as a duck as that she will see it as a rabbit).

As described, this mechanism implies that the 'aspect seeing' involved, whichever way it goes, is a product of perceptual learning. That is, it results from a simple associative mechanism requiring no involvement of beliefs, etc., and the effect, albeit ambiguous, might be assumed to be more or less permanent. If the phenomenon exemplifies cognitive penetration, it is therefore not in the phenomenon of aspect seeing itself, but in the subject's ability to 'flip' the phenomenology more or less at will. Here, it might be said that the subject's consciously entertaining the thought that the image represents a duck is what 'triggers' the switch to seeing it as a duck. Given my account of cognitive penetration, we might account for the aspect switch as resulting from the 'priming' of some potentially recollected objects more than others by means of some separate object of awareness. This separate, 'priming' object might be a label ('Duck') next to the ambiguous figure or it might itself be a recollected hearing of the spoken word 'duck', and so on. However the priming as achieved, it is liable to be described in terms of the subject's 'entertaining a thought about ducks' or some such, but it should be clear that no distinct 'thought' is required here beyond the recollection ('reawakened' awareness) of previously perceived ducks, and awareness of the priming stimulus (an utterance of 'duck', say).

Of course, aspect seeing or 'seeing-as' is a ubiquitous phenomenon not confined to ambiguous figures and the like. Most often, it is responsible for a certain veridicality in perception, namely our correct recognition of perceived objects as belonging to certain kinds. It is also, on occasion, responsible for error. We can account for error, too, using the simple metaphysical picture of perceptual learning proposed here. Consider Fodor's example of mistaking a cow for a horse (Fodor 1989: 107). To put it in Fish's terms, we here apply the wrong conceptual-recognitional capacity to the experience of the cow (2009: 165 ff.). However, the simple metaphysical picture explains the error in terms that acknowledge a genuine influence on phenomenology from the application of the wrong concept. What occurs in the case of seeing the cow is that, perhaps because of peculiarities of lighting or perspective, the cow causes activation of a perceptual neural state that is relevantly similar to that normally caused by horses. The occurrence and form of this neural state thus comes under the causal constraint of previously seen horses which thus help to shape the phenomenology of the experience.

If we ask what makes it 'wrong' for a sighting of a cow to elicit the re-presentation of previously seen horses then one answer is straightforward, not to say glib: *a cow is not a horse*. The model of concepts and perceptual content involved here is rather different to, say, Fodor's account, in which concepts are internal states or entities that represent horses or cows by virtue of some causal covariation with sightings of horses or cows. Beyond this causal connection, the link between representational state and what is represented is arbitrary, and it is this feature that generates a problem that Fodor's asymmetric dependence claim is intended to solve, *viz.* the problem that if causal covariation suffices for representation then the HORSE concept represents not just horses but also cows in bad light, large dogs in the mist, and so on (see Fodor 1989: 100 ff.).

Since the simple metaphysical picture does not posit inner representations this problem does not arise. On this view, what we call 'concepts' are better thought of as complex regularities of causal interaction between physical subjects and worldly objects. What do the work in perceptual experience – what contribute constitutively to its phenomenology – are not inner representations (or their representata) but worldly objects and their qualities.

Of course, more needs to be said about what makes a case of this sort an instance of misidentification and not just an instance of a cow putting us in mind of a horse. It should, for example, be a case in which the evocation of previously seen horses primes us to behave as if the thing seen really were a horse, e.g. to assent to "Is that a horse?", to put on a riding hat, etc. In other cases, there might be enough ambiguity in the cow's appearance to make us doubt our identification. In still other cases, we might recognise on balance that it is a cow and yet simultaneously enjoy the slight ambiguity and aspect shift that it engenders.

5.7 Conclusion - perception as presentation and re-presentation

I have proposed an account of the interface between the perceptual presence of our environment and our conceptual-recognitional capacities. This account combines the simple metaphysical picture of perceptual presence with its similarly object-involving account of memory and hallucination in order to build up a picture of the 'conceptual presence' of previously perceived objects in our current perceptual experience. Especially if such cognitive influences on phenomenology are widespread or even ubiquitous in perception, we might think that they undermine the natural view that perception is just a straightforward encounter with our environment. At least, they force us to acknowledge that perception is not always simply an encounter with what is before us right now. Even so, the account offered here still explains the cognitive aspects of phenomenology exhaustively by appeal to objects 'out there' in our environment, and without recourse to some 'internal' factors.

In summary, the account rests on a mechanism for perceptual learning (framed in terms of vision) which goes as follows:

- 1. The object O is before the subject S, and S sees O.
- 2. S's seeing O involves O's causing some pattern of neural activity N in S. N's occurring *now* is therefore a causal consequence of O's presence.

- Visually similar objects (seen in the same conditions) cause correspondingly similar patterns of brain activity in a given subject, both in respect of the location and form (pattern) of neuronal activity.
- 4. Normally, when an object perceptually causes a certain pattern of brain activity, the affected areas of the brain are altered such that their responsiveness to subsequent inputs of the same kind is changed in some way. The changes in part constitute what we call 'memory traces'.
- 5. That being so, and assuming a history of the subject's seeing objects like O, N's having the form that it does is a causal consequence not just of O's presence now but also of S's having previously seen relevantly similar objects.
- 6. Given the simple metaphysical picture, since the (visually perceptible) formal properties of the previously seen objects are partly causally responsible for N's having the form or pattern it does, those formal properties (and therefore the previously-seen objects) contribute to the phenomenology of the experience.
- 7. The visual experience of O has a character constituted not just by O and its properties but also by the visually similar previously seen objects.
- 8. Seeing O therefore 'brings to mind' in a rather literal sense those previously seen objects that belong to the same kind. Whether or not the subject is aware of the presence of those objects (recalls them as such), they exert an influence over the character of the experience, as revealed in cases of cognitive penetration.

I suggested that, with modest elaboration, the same mechanism could account for at least some of what is called cognitive penetration of perceptual experience. The extra step would be allowing an additional 'reawakening' of causal links to previously perceived Olike objects by way of associative links to other perceptible objects (labels, utterances, etc.). Much more needs to be said about the associative mechanism at play here to make a strong case for a simple metaphysical picture of cognitive penetration, and that is beyond the scope of the current thesis. It would likewise require a much larger project to assess whether the proposed mechanism for perceptual learning could form the basis for a full-blown theory of concepts. However, it does seek to marry the natural view with the claim that a conceptual-recognitional capacity serves at least as a bridge or interface between the perceptual (what we perceive in the sense of tracking) and the cognitive. As such, perceptual 'concepts' as I have described them appear as a certain kind of coming together or amalgamation of relevantly similar (current and previously) perceived objects in awareness, in such a way that they make a joint contribution to the overall phenomenology.⁵⁹ In light of the simple metaphysical picture, this 'coming together' of currently and previously objects is recast as a complex causal interaction mediated by the subject, and specifically her relevant neural state, a state whose occurrence and form is the causal upshot of all of those objects that contribute constitutively to the phenomenology of the experience. Thus, individual objects and the perception of them are clearly prior to the 'concepts' that are composed from them.

The proposed account also permits a partial rapprochement with the concerns of those philosophers who hold that experience is representational rather than merely presentational. Although the object perceived is not represented in experience, but merely presented, there is a sense in which perceptual experience does typically involve some representation. The 'conceptual-recognitional' element in experience is, in a rather literal sense, the *re-presentation* of objects previously perceived, objects that the currently perceived object 'brings to mind' in the manner described above. The explanation of the sensory and conceptual elements is thus unified, such that what is now a conceptual element was once a sensory element. Moreover, the proposed account makes it plausible to suppose that – except perhaps in some state of newborn inexperience – all of our perceptual experiences are 'conceptual' in the sense of involving such re-presentation to some extent. This would follow from a substantial overlap of neural loci for the effects of different experiences, as described above. The upshot is an explanation for both the immediacy and generality of perception cast simply in terms of awareness of worldly objects.

⁵⁹ It is an interesting question whether we should therefore say that the phenomenal character is borne by the scene before the subject. Our answer will hinge on whether we take the scene to be just what is physically most present to the subject, or the totality of what the subject is aware of. If the latter, then the scene is a complex amalgam of present and past objects of awareness.

CHAPTER 6 - PHENOMENOLOGICAL PROBLEMS (IV): THE SUBJECT-RELATIVITY OF PHENOMENAL QUALITIES

6.1 Introduction

In previous chapters I have described a number of perceptual phenomena that are ostensibly at odds with the natural view, and argued that they can in fact be squared with the natural view in one way or another. Most straightforwardly, many cases of perceptual illusion (notably what Fish calls physical illusions) can be accommodated to the natural view simply by noting that the latter does not require that worldly objects always appear perceptually to be how they really are. Rather, it requires only that the phenomenology of perception is explicable by reference to those objects and their qualities (plus background conditions of illumination, etc. and the way the objects are arrayed before the subject). This is consistent with the possibility that the way in which those objects and qualities are presented to us – perhaps in strange lighting or from an unusual angle – leads us to an erroneous judgment about the nature of the objects and qualities. Hallucinations, meanwhile, require a more sophisticated explanation, and here I appealed to the simple metaphysical picture as a way to account for the genuine phenomenology of hallucinatory experiences while avoiding the postulation of peculiar mental objects or qualities of awareness. The proposed solution - which takes hallucinations to have worldly objects furthermore offered a way to account for some more pervasive, everyday aspects of both illusory and veridical perception, namely the constitutive role in phenomenology of cognitive states such as memory, beliefs and judgments.

Common to these explanations was an attempt to explain the relevant experiences without compromising the natural view and, crucially, to do so while preserving the explanatory virtue of naïve realism. Thus, in each case I sought to account for the experience in such a way that its phenomenology is explicable by reference to worldly objects of awareness and their perceptible qualities. In other words, the order of explanation goes from the antecedent nature of the worldly objects and qualities perceived to the phenomenology of the experience.

Nonetheless, one or two cases were described where this approach appears somewhat strained. Key amongst these is the Rubik's cube illusion described in chapter 3.5, in which certain parts of the adjacent images, when seen in isolation from their surroundings, can be

seen to have the same colour, namely grey, but appear to have quite opposed colours – blue and yellow – when seen within the context of the images as a whole. Here, there seems to be a mismatch between how the image is, objectively, and how it subjectively appears to be. Most importantly, this mismatch seems not to be explicable solely by appeal to how things are in the subject's environment relative to the subject's spatial perspective, but instead appears to derive from the nature of processes internal to the subject.

In this chapter I will say more about this case as well as two further features of perceptual experience – in particular, of colour experience – that similarly point to what I will call the *subject-relativity* of phenomenal qualities. I will concede that these cases provide compelling evidence that, at least in the case of colour, the nature of the phenomenal qualities we perceive objects to possess is dependent on the nature of the perceiving subject. Put this way, the claim is ambiguous between several readings. On one reading, it might be taken merely to acknowledge that the nature of the subject – the scope of her sensory apparatus – is such that she is able to perceive only some worldly qualities and not others. That is entirely plausible but also perfectly compatible with the natural view, and as such is not relevant to my argument in this chapter. Rather, the interesting readings are those which claim that, *for some qualities actually perceived*, the nature of those qualities is determined in part by the nature of the subject.

As a first approximation, let us state the claim of subject-relativity as follows, formulated for the case of colour experience:

<u>Subject relativity</u>: the nature of perceived (phenomenal) colour is determined in part by the nature of the perceiving subject.

Further disambiguation is possible here. On the one hand, we might claim that when a subject looks at a coloured object, which colour she perceives it to have is dependent in part on the nature of that subject. On the other hand, we might claim that the nature of perceived colour *in general* is subject-relative. The nature of colour in general I take to be exhausted by the full range of visible hues together with their relations of mutual similarity (these hues and relations can be visualised in the form of a phenomenal colour 'space' as I will describe further in section 6.4 below).

In respect of the first reading, we might imagine two normal sighted subjects viewing the same object under the same conditions of lighting, perspective and so on. One sees the object as green, say, and the other sees it as blue (this is exactly the sort of case I will describe below as evidence for subject-relativity). Here, the relevant phenomenal quality of each subject's experience is relative to his or herself. This does not entail, however, that the nature of colour *in general* is subject-relative, just that which colour an object is perceived to have on any given occasion is subject-relative. We might very well claim in addition that the nature of colour in general *is* subject-relative, and this gives us our second reading of the subject-relativity claim.

Thus, again for the example of colour, we can state the two readings as follows:

<u>Subject relativity particular</u>: the colour an object is perceived to have is determined in part by the nature of the perceiving subject.

<u>Subject relativity general</u>: the nature of colour itself, including the full range of possible hues and their relations of mutual similarity, is determined in part by the nature of the perceiving subject.

I will cite some empirical findings concerning perception which compel us to concede the truth of both subject relativity ^{particular} and subject relativity ^{general}. Subject relativity ^{particular} is defended by appeal to two phenomena of normal colour experience: first, cases of perceptual colour adaptation (exemplified by the Rubik's cube illusion); and, second, marked intersubjective variation in hue perception. Subject relativity ^{general}, meanwhile, is supported by compelling evidence that the nature of colour in general is indeed subject-relative – specifically that the set of mutual similarity relations among colours, and which are partly constitutive of the nature of the colours, is explicable not by reference to properties of coloured objects but by the nature and functioning of the human visual system.

As it turns out, the use of two different kinds of case to defend subject relativity ^{particular} exposes some residual ambiguity in that thesis. In perceptual colour adaptation the subject-relativity of perceived (phenomenal) colour is illustrated by a phenomenon that is assumed to manifest similarly amongst all normally sighted subjects. Meanwhile,

intersubjective variation in hue perception illustrates the subject-relativity of colour by showing that the colour an object is perceived to have depends on which particular subject is doing the looking. In other words, perceptual colour adaptation shows that there is some feature common to normally sighted subjects – something in the structure and functioning of the normal human visual system – that partly determines and explains the (broad) colours objects are perceived to have in certain circumstances, while intersubjective variation in hue perception shows that differences between individual normally sighted subjects partly determine and explain more or less fine differences in the colours objects are perceived to have respective subjects.

These findings are of course not in conflict, and indeed may be mutually supportive. The fact that the specific natures of individual subjects partly determine and explain the colours they perceive objects to have is perfectly compatible with a broad commonality amongst subjects in respect of perceived colour. This follows straightforwardly if we factor in a broad similarity among normally sighted subjects in respect of the structure and functioning of their visual systems.⁶⁰

As such, there is compatibility between two kinds of phenomena illustrative of subject relativity ^{particular}. Although one could endorse subject relativity ^{particular} without endorsing subject relativity ^{general}, it is doubtful whether one could endorse subject relativity ^{general} without also endorsing subject relativity ^{particular}. If the very nature of colour is subject-relative then whatever colour an object is perceived to have on a particular occasion must also depend upon the nature of the subject. Note that this would not rule out the possibility that the nature of colour is determined by the nature of perceiving subjects but that all subjects nonetheless perceive the same objects as having the same colours. The subject-relativity of colour *per se* does not entail any intersubjective variation in colour perception, since all subjects might have visual systems that function in exactly the same way. Rather,

⁶⁰ The same point may be made to defuse a similar ambiguity that might be discerned in subject relativity ^{general}. So, on the one hand, we might take perceived colour to have a nature – revealed in the structure of similarity relations among individual colours – that is common to all normally sighted subjects because determined by some features of the visual system shared by all these subjects. On the other hand, we might find that the nature of colour is not common to all normally sighted subjects but varies according to the precise nature of the subject (the structure and functioning of her visual system). In this latter case, the nature of colour in general – in particular, the structure of relations of mutual similarity among perceived colours – might vary from subject to subject, such that there is no single colour nature shared by all normally sighted subjects. Conversely, were such intersubjective variation lacking, this would be explicable in terms of the identical functioning of different subjects' visual systems.

it entails that the perceived colour of an object is subject-relative and cannot be explained without reference to something in the nature of the subject.

Although subject relativity ^{particular} is thus consistent with the denial of subject relativity ^{general}, we might hope for a unified explanation of the two theses if it turns out that there is good independent evidence for both. The phenomena of colour vision that I will describe in this chapter provide just such evidence, and I will go on to offer an account, employing the simple metaphysical picture, that might point to such a unified explanation.

In the meantime, if one or both varieties of subject relativity are true then the natural view looks to be in trouble. After all, if some or all phenomenal qualities are subject-relative then surely they cannot be out there, inhering in worldly objects as they seem. That would be too hasty a conclusion, however. What we are entitled to conclude from the conjunction of the natural view and subject relativity is that phenomenal qualities are out there in the environment, and that these qualities are subject-relative. However absurd that conclusion might seem, it is at least consistent and intelligible. The seeming absurdity reflects the fact that conjoining the natural view and subject relativity yields the denial of the realism claim, namely the claim that worldly objects exist and have their perceptible qualities independently of their being perceived. Since I defined naïve realism as the conjunction of the natural view and the realism claim, it follows that the truth of subject relativity entails the falsity of naïve realism but not the falsity of the natural view.

Still, abandoning the realism claim would seem a very high price to pay for retaining the natural view. It would leave us with two major questions:

- 1. How can we make sense of qualities' being out there in the environment and yet also determined in part by the subject?
- 2. What does this entail for other qualities of objects that do not seem to be subject-relative, and indeed for the objects themselves?

The second question reflects a seeming disparity thrown up by subject relativity between those qualities that seem subject-relative (colour, in the cases to be described) and others that seem quintessentially subject-independent, such as shape and mass. To suggest this distinction is of course to echo the distinction between primary and secondary qualities proposed by Galileo, Locke and others. Those thinkers were led by this distinction to conclude that certain qualities are really out there (primary qualities) while others exist only 'in the mind' (secondary qualities). The natural view plainly requires that we reject the latter confinement of colours, etc., to the mind, but accepting something like the primary/secondary quality distinction forces us then to conclude that, of the perceptible qualities of worldly objects, some are subjective (subject-relative) and some are objective (subject-independent). This is worrying since it is rather implicit in the natural view that all perceptible qualities are 'on a par': when a subject looks at a tomato it is the tomato's roundness and the tomato's redness that together constitute the relevant part of the experience's phenomenology. According to the natural view we are aware of these (respectively, 'primary' and 'secondary') qualities in the same way – simply by their being out there in the world and our seeing them. What's more, the roundness and redness in this instance are rather literally bound together – the redness pervades the tomato's outer surface which in turn defines its shape.

Any solution to the second question is likely to hinge on our answer to the first. As it turns out, the view that certain qualities are out there in the world but nonetheless subject-relative (or 'mind-dependent') can in fact be found quite widely among philosophers, albeit not allied to anything like the natural view. The kinds of perceptual phenomena motivating subject relativity are often used to support the broad kind of theory known as *relationalism* about colour (for a defence see e.g. Cohen 2004, Cohen 2007). Several varieties of colour relationalism have been proposed (for a brief summary see Cohen 2010: 232), but one in particular has been widely defended and conforms, at least on the face of it, with the view (explicit in the natural view) that colours are out there inhering in the ostensibly coloured objects perceived. This kind of colour relationalism holds that colours are dispositional in nature – specifically, an object's colour is its disposition to look a certain way (red, green, etc.). So, for a tomato to be red is for it to be disposed to look red to normal subjects under normal conditions (of illumination etc.). The redness here is identified with the disposition to look red, and is therefore subject-relative since its being the quality it is tied essentially to how it looks to normal subjects.⁶¹

⁶¹ There is a related account of colour which says that what it is for an object to be red is for it *to have a property* that disposes it to look red to normal subjects under normal conditions. Where the relevant property is taken to be physical, this can be considered a form of physicalism about colour (see e.g. Glüer 2007).

I will later remark on the unsuitability of a dispositional account of colours for a theory endorsing the natural view. Unfortunately, this will leave few, if any, obvious alternatives. Several of the more promising theories of colour will prove casualties of the perceptual phenomena I am about to describe. Notably, these theories include those that identify them with physical properties of objects (one or other form of physicalism), and those that, by contrast, treat them as occurrent properties of objects that are primitive, irreducible to physical properties (this is primitivism about colour).

Mindful of some of the features of colour experience that I am about to describe, Johnston claims that

> when it comes to the external explanatory causes of our color experiences, psychophysics has narrowed down the options. Those causes are either nondispositional microphysical properties, light dispositions (reflectance or Edwin Land's designator dispositions or something of that sort) or psychological dispositions (dispositions to appear colored) with microphysical or light-dispositional bases.

> > Johnston (1992: 224)

It will transpire that neither the physicalist nor the dispositionalist theories of colour are a good fit for my defence of the natural view, since they either fail to explain certain features of colour experience or fail to accord with the basic insight of the natural view, *viz.* that phenomenal qualities like colours are out there in the environment, inhering in worldly objects. I will argue, however, that if Johnston's claim therefore seems to leave us without any remaining options, then that is in part because he is constrained by the assumption that experiences are caused by their objects. A rejection of this assumption, via endorsement of the simple metaphysical picture, points the way to an alternative account of colours, one that acknowledges their subject-relativity while leaving them just as much 'out there' in the world as the shape and size of things, and indeed the things themselves. First let us turn to the allegedly problematic features of colour perception.

6.2 Colour adaptation phenomena

A striking feature of colour perception is that the apparent colour of objects is susceptible to variation according to the visual context in which they are seen. In chapter 2

I described the phenomena of simultaneous colour contrast and colour constancy, in which an object's apparent colour is influenced, respectively, by the colour of its surrounds or by the ambient illumination. Both phenomena can be understood as the upshot of compensatory processes in the visual system that enable subjects to track unchanging surface properties of objects through changes in visual context.⁶² For that reason I will collectively call these *colour adaptation phenomena*.

In the normal course of events, colour adaptation actually serves to *limit* changes in the apparent colour of objects despite changes in visual context. So, for example, surfaces often appear to retain their uniform colouration despite spatial or temporal variations in illumination. However, I described cases in which the same phenomena give rise to illusion, and in these cases the context effects clearly alter the colour of the focal object or stimulus. For example, in the Rubik's cube illusion the two images are presented as though depicting the same or similar object situated in different background lighting conditions (or as seen through differently-coloured glasses). The left-hand image appears to show the cube as if it were seen in a yellow light or through a yellow filter, while the right-hand image appears to show it in a blue light or through a blue filter. One interpretation of what occurs in viewing these images is that we are led to judge mistakenly the colour of certain of the squares depicted as making up the surface of the cubes: certain of these squares are judged to be blue in the left-hand image and yellow in the right-hand image, when in both they prove to be grey when viewed in isolation. However, the peculiar force of the illusion is precisely that it seems inadequate to explain it in terms of judgment, at least insofar as we are to draw a distinction between judgment and awareness in the way proposed by the two-faculty view discussed (and dismissed) when considering illusory perception in chapter 3. What makes the illusion so powerful is that it seems to insist that the seeming blueness or yellowness of the squares in question are not merely matters of judgment, but are genuine qualities perceived – genuine phenomenal qualities.⁶³

⁶² Similar perceptual constancy is observed in respect of other perceptible qualities such as shape and size (as when objects are seen from different perspectives).

⁶³ In this respect, the illusory cases avoid or dissolve an ambiguity that seems to attend the normal (loosely 'veridical') upshots of colour adaptation. For the philosopher of perception these latter cases can give rise to difficulties in disentangling different senses of 'looks': should we say that the white car under the sodium streetlamps looks both white (perhaps in the epistemic sense of 'looks') and orange (in the phenomenal sense)? The more striking illusions, such as the Rubik's cube illusion, are striking (i.e. clearly illusory) precisely because such ambiguity seems absent – the target squares in no way seem grey, the colour that they are seen to possess when viewed in isolation.

On the face of it, colour adaptation effects might seem to be explicable wholly in terms of how things are out in the world. The phenomena seem to be universal and predictable, with the perceived colour of an object being systematically related to such 'external' features as the object's 'actual colour' (how it would look in optimal viewing conditions, however these are to be specified), its size and orientation relative to the subject, the colour and dimensions of surrounding objects, ambient illumination, and so on. Colour adaptation effects would therefore seem to exemplify what Fish calls *physical* illusions (noting of course that where the same phenomenon underwrites the appearance of colour constancy we would not attribute any illusion). However, at least some of these contributory features would not seem sufficient genuinely to explain the perceived colour of the target object. We can indeed understand how ambient illumination can change the apparent colour of an object (the white cube in red light, for example), namely by determining the actual spectral profile of light reflected from the object (to something like the profile that would be reflected by a red cube in white light). By contrast, colour adaptation effects need not result from any change in the object itself or in the nature of the light it reflects. To understand them fully it seems that we need to look elsewhere than in the subject's environment – specifically, in some compensatory mechanisms in subjects' visual systems that modulates the responses to localised colour stimuli according to the colour of background objects and/or the colour of ambient illumination.⁶⁴

Colour adaptation illusions thus support subject relativity ^{particular}, namely the claim that the perceived colour of a given object depends in part on the nature of perceiving subjects – in particular, the structure and functioning of their visual systems. Nonetheless, we might envisage one possible explanation of this that renders it harmlessly compatible with both the natural view and the realism claim. According to this explanation, the altered colour appearance of the object when seen against differently-coloured backgrounds might be due to changes in subjects' sensitivity to certain components of the object's colour due to adaptation effects (photoreceptor fatigue) arising from the background colour. For example, subjects viewing a purple object against a red background might become less sensitive to the red component of the object's colour,

⁶⁴ For the classic explanation of simultaneous colour contrast in terms of opponent-process theory, see Jameson & Hurvich (1959), and for an alternative theory in light of more recent contradictory evidence see Ekroll & Faul (2012). Although simultaneous contrast effects contribute towards colour constancy, the latter relies on compensatory responses to diverse confounding variables (e.g. orientation, depth, specular highlights and patterns of light and shade, in addition to changes in the colour of ambient light) so it seems likely that more than one physiological mechanism is involved.

effectively emphasising its blue component instead. To this extent the phenomenon might be explicable in terms of effects at the photoreceptor level, and the change in perceived colour might be attributable to partial insensitivity to an object's existing colour rather than to awareness of a *different* colour. This explanation would be compatible with the conjunction of the natural view and the realism claim (i.e. naïve realism) since it would leave unhindered the claim that the relevant phenomenal quality (blue) is just the blueness mind-independently inhering in the object perceived. The failure of the red component of the object's purpleness to feature in phenomenology would simply reflect a (temporary) limitation on the subject's perceptual sensitivity, where such limitations (permanent in the case of ultra-violet, for example) are perfectly compatible with the natural view.

Certainly, the established physiological theories of simultaneous colour contrast allow that the effect may arise in part from adaptation to preceding and background colour stimuli in just the way described. However, they also emphasise the more important influence of downstream processes in the visual system (e.g. Ekroll & Faul 2012). In any case, an appeal to partial sensitivity would not account for colour adaptation effects experienced when looking at objects that are, under more optimal conditions, seen as possessing one of the unique hues, for in such cases the object's colour has only a single colour component. More strikingly still, partial sensitivity cannot account for the sort of colour adaptation effect exemplified by the Rubik's cube illusion, where the target object is grey when viewed in isolation (i.e. altogether lacking in hue) and yet appears variously blue and yellow depending on its background.⁶⁵

⁶⁵ We might try to explain the Rubik's cube illusion as a form of perceptual learning. Thus, we might suggest that the appearance of blue or yellow where there is 'objectively' an instance of greyness is due to a learned grasp of how colour appearance varies under different lighting conditions. Accordingly, when something is presented to us as achromatic under blue illumination (as depicted in one of the illusory images) perhaps we associate this with its being 'objectively' yellow (i.e. yellow in broad daylight) and it is the yellowness that is thus 'brought to mind' which influences the phenomenology of the experience. However, some sort of associative process appears unlikely here. Notably, the colour adaptation phenomenon that seems chiefly responsible for the Rubik's cube illusion is also observed when subjects view images that are not intended to depict the sort of real-world scenarios that might be expected to elicit learned associations. Most typically, simultaneous colour contrast is illustrated with quite abstract images of a small square of one colour centred within a larger square of contrasting colour. Although simultaneous contrast effects are typically hypothesised to reflect an evolved compensation of the visual system to variations in background illumination (Lotto & Purves 2000; although see Ekroll & Faul 2012 for an alternative explanation), these sorts of abstract image do not *look as though* they involve or depict an instance of something's being seen in coloured light. The illusion therefore does not seem to involve or require any recognition or conceptualisation of the image as a depiction of something (a small square) seen in coloured light. This tallies with evidence that the compensatory mechanism involved in simultaneous colour contrast occupies a stage in early visual processing, largely at the level of the opponent process (Ekroll & Faul 2012).

6.3 Intersubjective variation in unique hue perception

While colour adaptation illusions seem to exemplify a form of subject-relativity that is universal amongst normal subjects, I will describe in this section a further phenomenon of colour perception whose interesting feature is that it presents as a *difference* between the colour perceptions of different subjects. This therefore supports subject relativity ^{particular} via a somewhat different but complementary route.

It is widely held that a universal and essential feature of perceived colours is their intuitive division into those that are 'pure' or 'unique' and those that appear to be formed from a mixture of those unique colours (see e.g. Hardin 1997: 291-2; Kay & McDaniel 1997: 411; Conway & Stoughton 2008). Thus, unique green is a green which appears neither bluish nor yellowish, and unique blue is a blue that appears to have a trace of neither green nor indigo. Purple, however, is a 'mixed' or 'impure' colour that looks to be formed from blue and red. Although different subjects typically agree that there is this distinction between unique and mixed colours (hues), there is often marked differences in the objects different subjects will point to as exemplars of the unique hues. It is this intersubjective variation in unique hue perception that seems to support the thesis subject relativity particular. Although any given subject may be consistent in attributing, say, unique greenness to the same object or spectral light at different times (assuming viewing conditions remain the same), a different subject is likely to (again, consistently) attribute unique greenness to spectral light of different wavelength or an object of different spectral reflectance (Kuehni 2004; Webster et al. 2000).⁶⁶ Importantly, this is a phenomenon of *normal* colour vision; in other words, it is observed amongst subjects who have no recognised abnormality of colour vision, such as any of the various forms of 'colour blindness'. The intersubjective variation afflicts all of the unique hues (red, yellow, blue and green) but is especially marked for green (Kuehni 2004: 161). As Block notes, so marked is it in the case of unique green that, among a random sample of normal subjects, there is never likely to a be a majority consensus over the identity of the unique hue: "if we

⁶⁶ There is also evidence for some *intra*subjective variation in unique hue perception over longer periods of time. In particular, the wavelength of monochromatic light sources seen as unique green shorten as subjects age, and this seems likely to be due to changes in neuronal sensitivity at the level of the opponent process (Schefrin & Werner 1990). The same study also found that an individual subject will attribute unique green to light sources of slightly different wavelengths depending on the intensity (luminance) of that light source (ibid.).

take a chip that any one subject in this experiment takes as being unique green, most of the others will see it as at least slightly bluish or yellowish" (Block 1999: 43). More striking yet is the observation that there is overlap between the range of hues identified as uniquely blue and that identified as uniquely green – so there will be some hues that some subjects identify as uniquely green but others identify as uniquely blue (Kuehni 2004: 161).⁶⁷

The outline of the problem is as follows:

- 1. Two different normal subjects, given the same viewing conditions, see the same stimulus (object or spectral light) as being of a different hue.
- 2. The stimulus is presumed really to possess a single hue.
- 3. Given (1) and (2), one or both subjects must be misperceiving the stimulus.
- 4. However, there is no principled means of deciding which subject is misperceiving the stimulus.

It is not known for certain why such intersubjective variation in hue perception occurs, although it presumably reflects some variation in the structure or functioning of subjects' visual systems. There is known to be genetic variation in retinal cone pigments between subjects which affects the peak spectral sensitivity of those pigments, and this is likely to explain at least some of the variation (Lutze et al. 1990; Neitz & Jacobs 1990). If this is correct – if variation in colour perception is attributable to variation in the structure or function of subjects' visual system – then we have strong grounds for concluding that the internal physical make-up of the subject at least partly determines the phenomenology of perceptual experience. In other words, the phenomenon of intersubjective variation in hue perception appears to support subject relativity ^{particular} – the colour an object is perceived to have is determined (and explained) in part by the nature of the perceiving subject.

Intersubjective variation in unique hue perception has been presented as a challenge to any form of objectivism about colour, namely the view that colour is an objective, mind-

⁶⁷ As Allen warns, however, this finding results from comparing results across different studies, which leaves open the possibility that the apparent intersubjective overlap of unique green and unique blue might be an artefact of different experimental conditions (Allen 2010a).

independent feature of the world. Objectivism takes various forms (see Maund 2012, s.6.2 & 6.3) but, crucially for our purposes, it would include naïve realism of the sort that endorses the natural view. If we assume that (2) above is correct – that in a set of different hue samples there can be at most one that is unique green – then the naïve realist is committed to claiming that most of the normal subjects, even in optimal viewing conditions, misperceive the stimulus. Yet, as (4) points out, all of the subjects in the relevant experiments are considered 'normal', so we have no way of determining what counts as misperceiving in these cases.

It is easy to see how the problem of unique hues can lead to a denial of the mindindependence of colour. Certainly, the problem of unique hues generates considerable tension between the 'naïve' and 'realist' aspects of naïve realism: if we hold onto the claim that phenomenal properties are 'out there' in the physical environment of the subject, then this kind of intersubjective variation seems to entail that the properties are not mindindependent. As Kalderon puts it, "if the appearance of unique green depends on the visual system of the perceiver, then it is hard to understand how it could be the manifestation of a *mind-independent* quality" (2007: 589).

Naïve realists and representationalists alike have sought to defend their views in light of this phenomenon.⁶⁸ One response has been simply to insist that some subjects and not others accurately perceive the real hue of objects. For example, it might be that organisms have evolved to perceive accurately only *general* or *coarse-grained* colours (blue rather than Prussian blue, perhaps), but vary in their accuracy in perceiving the determinate, finegrained colours perceived (Allen 2010a; Tye 2006). The explanation for this might be that there has been no adaptive value in accurately perceiving fine-grained as opposed to coarse-grained colours (Tye 2006). However, this explanation faces several objections. For one thing, the proposed solution requires that some ostensibly normal perceivers and not others misperceive the (fine-grained) colours involved, and yet the phenomenon is problematic precisely because there is no non-arbitrary way of deciding which subjects perceive accurately and which do not (Cohen, Hardin & McLaughlin 2006: 338-9). In addition, and as noted above with the example of unique blue and unique green, subjects

⁶⁸ See e.g. Tye (2006) for a defence of representationalism in light of variation in unique hue perception. This is the kind of strong representationalism which is 'phenomenally externalist' in holding that phenomenal qualities are to be identified with the (worldly) qualities represented by the experience (e.g. Dretske 1995, Lycan 2001, Byrne & Tye 2006).

may differ in their colour experience even to the extent of assigning the same objects to different coarse-grained colour classes (although see footnote 67 above).

As it is, even if we think the appeal to coarse-grained colours suffices to uphold some form of direct realism (perhaps including Tye's representationalist sort), it does not appear a satisfactory solution for a proponent of the natural view, for whom the phenomenal quality in veridical colour experience is just the colour itself. Since we *can* discriminate fine-grained shades, it seems to follow from the natural view that the fine-grained colour attributions of normal subjects should agree under the same conditions of viewing.

Fish offers a solution to the problem of unique hues which seems to get around this worry: perhaps subjects who differ in their colour experience are, as a result of differences in the sensitivity or functioning of their respective visual systems, aware of slightly different components of the colours that are really out there (2009: 154, fn. 3).⁶⁹ However, Fish's solution looks untenable for two reasons. Firstly, unique hues are precisely those that lack multiple colour 'components', so it should not be possible for the 'real' unique hues to be seen as mixed by any normal subjects (when viewing in white light). As such, a proponent of the natural view will struggle to explain how an objectively pure green object could appear blueish to some subjects and yellowish to others, if there is in fact no blueishness or yellowishness out there. All normal subjects should therefore agree about the location of unique hues, and any disagreement should only pertain to nearby mixed hues whose minor component some 'insensitive' subjects are unable to perceive. But this is not what the evidence indicates: subjects often disagree entirely about the location of the unique hues, with no overlap in location.⁷⁰ It would follow from Fish's suggestion that there *are* no unique hues – that hues are always mixed but sometimes appear unique.

⁶⁹ Fish also suggests that the differences are not really in what phenomenal colour qualities are presented in their respective experiences, but are merely differences in the subjects' colour judgments or attribution of colour terminology (Fish 2009: 154, fn. 3). However, this effectively amounts to a denial of the supposedly reliable intuitive distinction between unique and mixed hues on which the whole claim is based.

⁷⁰ The response to Fish's first solution also undermines Kalderon's preferred solution, which is what he calls 'colour pluralism'. This is the claim that objects may have many colours, not all of which will be visible to any given subject (Kalderon 2007). Subjects' sensitivity to colour properties might, in this sense, be only partial. But, again, a restricted sensitivity to some hues should still leave agreement about those hues that are simple (unique), and this is contradicted by the evidence. Kalderon, indeed, admits that subjects are very consistent in attributing unique hues to narrow ranges of spectral light, whereas intersubjective variation is great (ibid.: 565).

The second counterargument to Fish's proposal appeals to the physical nature of the light stimulus. Subjects can perceive a unique hue when looking at an object in white (i.e. broad-spectrum) light, such that the object reflects light of various wavelengths. This might seem potentially to lend support to Fish's suggestion: perhaps different subjects are more or less sensitive to different parts of the visible spectrum of wavelengths. However, unique hue perception can also occur when subjects are shown sources of light of a single wavelength, and again there will be disagreement among subjects as to which wavelengths correspond to the unique hues. It therefore appears in this case that the physical stimulus itself embodies no mixture of components (no range of different wavelengths) that could underwrite Fish's appeal to varying sensitivities. If one subject is less sensitive than another to a certain wavelength of light, we might imagine that she will see it as a less intense colour, but we would expect her to see it as the *same* colour, unique or otherwise.

We should note that this second objection to Fish is not conclusive: one might counter that it relies too heavily on the assumption that there is a close relationship between the nature of light and colour. One might, for example, be a primitivist about colours and just take them to be *sui generis* properties of objects distinct from their physical properties.⁷¹ However, there are good empirical reasons to think that colour perception *is* closely tied to the nature of the light stimulus (but, crucially for my argument, also to the nature of the physical subject). As I argued in chapter 2, there are ample reasons to think that an appropriate causal link between subject and object is necessary for perception, and we might see the physical medium (light, in the case of vision) as an effective 'bottleneck' in terms of constraining what we can and cannot perceive of the object. If the mediating stimulus is physically simple and homogeneous (as with light of a single spectral wavelength) then it is unclear how it could mediate the awareness of something complex or heterogeneous. And yet, if the natural view is correct, it does precisely this. After all, although subjects will see some single-wavelength light sources as being of a unique hue, those sources of intermediate (but still single) wavelength will be perceived as binary hues.

If the various attempted explanations for the intersubjective variation in hue perception fail, then we are left with good evidence for subject relativity ^{particular}: which colours an object is perceived to have are subject-relative inasmuch as they can vary

⁷¹ For critical discussion of colour primitivism see Byrne & Hilbert 2007. I will discuss primitivism further below, rejecting it as a possible explanation for subject relativity general.

markedly from one subject to another. This complements the finding of the previous section, that colour adaptation effects reveal the subject-relativity of perceived colour in ways that are likely to be shared by *all* normal subjects. In the next section we can go further and consider evidence for subject relativity ^{general}, namely that the nature of colour itself is subject-relative and can only be explained by reference to the ways in which the subject responds physiologically to perceptual contact with coloured objects.

6.4 The nature of colour

In the previous two sections we have considered two kinds of variability in colour perception that demonstrate the subject relativity of phenomenal colour in at least the first sense (subject relativity ^{particular}). In the first case – colour adaptation – the phenomenon reveals *intrapersonal* variation in colour perception: an object viewed under the same lighting conditions but against differently coloured backdrops will look to any given subject to have a different colour in each case. The example of the Rubik's cube illusion was used to press home the conclusion that this difference in perceived colour was a genuinely phenomenal matter and not a matter of mere judgment. Crucially, although the variation in perceived (phenomenal) colour is shared by different subjects and appears to be consistently and systematically related to the layout of the scene perceived, it is not adequately *explained* by the latter. Rather, it seems better explained by the nature of adaptive visual mechanisms internal to the subject.

The second kind of phenomenon under consideration revealed in addition a marked *intersubjective* variability in colour perception, specifically in the way different subjects will assign unique hues amongst diverse coloured objects or light sources. In this section I will show that subject relativity of colour would still hold true even if there were no evidence for either intrasubjective or intersubjective variation. Even if all subjects saw the same things to be the same colours, absent any simultaneous contrast effect, we would still be compelled to the conclusion that perceived colour is subject relative. This is because the very nature of colour itself appears to be determined by the nature of subjects – specifically, the structure and functioning of their visual systems. This amounts to what I called subject relativity ^{general}.

To appreciate this we first need to establish where to look for the *nature* of colour. It is important to reiterate that it is perceived or phenomenal colour that we are interested in. Some theorists about colour will distinguish between the colours that objects seem to have (phenomenal or 'subjective' colours) and those that they 'really' have ('objective' colours), where the latter might be identified with physical properties of objects' surfaces, etc. (see e.g. Harman 1996). Phenomenal colours might, on this view, be identified with qualities of (inner) experiences (qualia), instances of which are caused to occur by the worldly properties that are identified as the real colours. In any case, the natural view is a claim about phenomenal qualities, and phenomenology overall, and it furthermore identifies these with qualities inherent in the worldly objects perceived. On this view, colours are as 'real' as other qualities of objects, such as their shape and size.

Now, in seeking the nature of colour, we might look for those of its properties that are considered essential to it. In this respect, there is widespread agreement that the essential properties of colours include not only how they look but also, and derivatively, how the various colours relate to one another in terms of similarity and dissimilarity. It is possible to illustrate these relations by plotting colours within a three-dimensional 'colour space' whose three axes correspond to the three principal attributes of colour, namely hue, saturation and brightness. If we disregard the dimension of brightness (which is less obviously subject-relative, being a function of the luminance or amount of light emitted or reflected from an object) we obtain a two-dimensional array of hue and saturation which together define objects' *chromaticity*. It is the structure of this chromaticity space that is most obviously subject-relative.

One reason for thinking so is that the structure of the chromaticity space is not correlated closely with any physical features of coloured objects. For example, given a spectral light source, one can fix all of its physical properties (e.g. amplitude, polarity) while varying only its wavelength such that the light is of a single wavelength at any one time but this wavelength varies over time. This light appears to vary in hue as its wavelength changes. By gradually progressing from light of wavelength 380nm to light of wavelength 750nm a range of ('spectral') hues is generated which include all the colours of the rainbow in sequence from violet to red respectively.⁷² Strikingly, although the physical

⁷² Such spectral hues do not exhaust all the visible hues, which include hues such as purple which are obtainable only by mixing light of different wavelengths.

characterisation of the stimulus admits only continuous variation in a single dimension (light wavelength), the corresponding hues exhibit mutual relations of similarity that can be adequately modelled only in four dimensions corresponding to the unique hues, which for normal humans are red, yellow, green and blue. Implicit in this multi-dimensionality of hue is the distinction, exploited in the argument for intersubjective variation in colour perception, between unique and mixed hues. Although there are specific spectral hues that subjects will pick out as 'pure' or unique, there is nothing physically distinctive about the wavelengths of light that correspond to unique hues as compared with those that correspond to 'impure' or mixed hues. Thus, while we can relate changes in spectral hues to changes in light wavelength, there is nothing in the nature of the latter that explains the nature of the former.

What's more, the relationship between phenomenal hue and the properties of emitted or reflected light is even more complicated than the example of spectral hues suggests. Objects that emit or reflect visible light of different wavelengths can nonetheless exhibit the same colour (hue, saturation and brightness) to a single subject, a phenomenon known as *metamerism* (see e.g. Hardin 1988: 28). Metamerism does not occur between sources of spectral light, but only between objects (metamers) at least one of which reflects or emits light of various mixed wavelengths. Such metamers will exhibit the same colour only under certain illuminations and not others (assuming they reflect rather than emit light).

Taken together, the nature of spectral hues (including their similarity relations and the unique/mixed hue distinction) and the phenomenon of metamerism show that there is no clear correlation between phenomenal colour and the physical properties of the light that we take to mediate colour experience.⁷³ In contrast, it has long been recognised that the

⁷³ It is worth acknowledging that a focus on light risks misplacing the locus of perceived colour, at least in respect of those ordinary objects we see in virtue of their reflecting light. Here, the colour is typically perceived as inhering in the object's surface (or suffusing the object if translucent and transmitting as well as reflecting light); arguably, we do not take ourselves to be seeing the light at all in these cases, except where the object sports some specular highlights. Leaving aside the fact that we do see sources of light as coloured, and that these exhibit the same colour properties including metamerism, shifting our attention to the physical surface properties of reflective objects only makes matters worse. Identifying colours with intrinsic properties of objects' surfaces, whether structural, chemical or microphysical, has proved even less satisfactory. Physically diverse surfaces can appear to have the same colour, and these surfaces need not even be metamers – they might be isomers, possessing the same overall reflectance properties that some philosophers have claimed that colours are instead best identified with the propensity of those surfaces to reflect light with certain wavelength profiles and intensity; in other words, an object's ('objective') colour is identified with its *surface spectral reflectance* (ibid.: 63-4). This move, of course, leaves us with the same problems arising from our attempt to correlate perceived colour with the physical properties of light.

structure of phenomenal colour space is much better explained by the structure and functioning of the human visual system. In other words, phenomenal colours are poorly correlated with the physical properties of the light stimulus but highly correlated with the way that we physiologically react to that stimulus. One early theory – trichromatic colour theory – suggested that the structure of perceptual colour space arises because the human eye does not respond in a linear fashion to changes in light wavelength; instead, it was proposed that there were three different kinds of receptive component in the eye, each of which was sensitive to a different range of light wavelength and would yield in response a different pure colour experience of either blue, green or red. It was the relative activation of these different receptors that was taken to determine the colour perceived by the subject. The mechanism would be additive: equal activation of the 'blue' and 'green' receptors would, for example, yield an experience of turquoise (Hardin 1988: 29).

Although it was subsequently proved that the human eye does indeed contain three kinds of photoreceptor (retinal cone cells) sensitive to light of different wavelengths, the simple model proposed by trichromatic colour theory has proved inadequate to explain the structure of phenomenal colour space. In particular, it was noted by early critics of trichromatic colour theory that certain colour combinations that we might expect from the model do not in fact occur. For example, a mixture of lights that individually look red and green results in the appearance of yellow, and not a greenish-red as we would expect from the theory; indeed, certain expected colour combinations, notably greenish-red and blueish-yellow simply do not occur in human colour space (Hardin 1988: 29). Since there are indeed three kinds of 'colour receptors' in the eye, each responding maximally to different regions of the visible spectrum, the trichromatic theory would seem to hold at least as far as the retina. As it turns out, however, the peak sensitivities of these colour receptors do not in fact perfectly match wavelengths of light that we see as pure green, red and blue, the mismatch being especially great for red.

Thus, to account for the seemingly anomalous features of colour space we must to some other features of visual processing occurring downstream of the retina. The impossibility of certain colour pairs (blue-yellow, red-green and black-white) suggested that these colour pairs correspond to processes in the visual system that stand in some form of opposition, i.e. activation of the 'blue' process in some way inhibits or rules out activation of the 'yellow' process. This theory was therefore called the *opponent process* *theory*, and it has received broad empirical support, even if it is not in its current form a perfect fit for the various lines of evidence from colour science (Hardin 1988: 35). The appearance of the various colours thus corresponds to post-receptoral processes – those downstream of the retina – whose structure results from interactions between post-receptoral neuron groups whose activity is not merely additive but also mutually inhibitory.

It seems then that the best explanation of the correlation between the structure and function of the human visual system and the structure of perceptual colour space is that the former somehow determines the latter. In other words, we might infer from the observed correlation the truth of subject relativity ^{general}: the nature of colour itself, including the full range of possible hues and their relations of mutual similarity, is determined in part by the nature of the perceiving subject.

As Pautz puts it, "even under ideal conditions (no interfering factors), sensory character is much better correlated with neural patterns in the brain than with anything in the external world" (2014: 238). If we present a subject with a range of objects of, say, different colours, their subjective reports of relative chromatic sameness and difference will prove to correlate more closely with differences in the neural goings-on in her brain than with physically describable properties of the objects themselves (their surface spectral reflectance perhaps). This presents what Pautz calls a 'problem of correlations' for those, like naïve realists, who wish to identify phenomenal qualities with qualities inhering in worldly objects (ibid.).⁷⁴

Pautz marshals a good deal of compelling evidence to show that the problem of correlations arises not just in respect of visual qualities (colours) but also in respect of qualities accessible to other sense modalities, including taste, smell, sound and pain.⁷⁵ Here again, the apparent (phenomenal) relations of similarity and difference among

⁷⁴ Pautz's key target in his (2014) is what he calls 'tracking intentionalism', namely any phenomenally externalist form of representationalism which holds that phenomenology is determined by whatever one represents visually in having an experience, where a subject's representing in this way involves some sort of 'tracking' or covariation relationship between some internal (neural) state of the subject and the property represented (Pautz 2014:: 240). As he notes, however, the argument generalises to all forms of phenomenal externalism, including naïve realism (ibid.: 238).

 ⁷⁵ As Pautz points out, although there is 'bad external correlation' for qualities like colour and smell, there are (fewer) qualities for which there is better external correlation, including shape and size (Pautz 2014: 242). This functional difference among qualities is clearly aligned with the classic 'primary/secondary quality' distinction. I will discuss this further in chapter 7.

qualities prove to be better correlated with patterns of neural activity in the perceiving subject than with physical properties of the external stimuli or the objects in which those qualities seem to inhere. For example, some substances of closely similar chemical structure can smell quite markedly different, while other substances of very *different* chemical structure can smell more or less the same (Pautz 2014: 246-7 & 269-71). And the explanation for this appears to lie in the extent to which the different chemicals elicit the same or different patterns of neural activity in a key part of the brain involved in olfactory perception (ibid.).

Pautz takes these findings concerning the nature of perceived qualities to reveal what he calls their *internal dependence* (2014: 252 ff.; see also Pautz 2011: 402-5). This is clearly in tension with the notion that colours are mind-independent features of the world – what Pautz calls their *response independence* (2014: 239; 2011: 388). Given internal dependence, one might be tempted to conclude that colours, etc., are in some sense 'in the mind' or 'in the head', as in sense-datum theories or the sort of internalist identity theory which holds that types of phenomenal qualities are identical to types of brain states. Pautz, meanwhile, opts for an 'Edenic' theory of the sort proposed by Chalmers, according to which phenomenal qualities are not real qualities at all truly to be located *anywhere* – they are instead "wholly chimerical" (Pautz 2014: 295; see also Chalmers 2006).

Is there a way to resist such conclusions and defend the natural view? One option might be to endorse *primitivism* about colour: the view that colours are primitive, *sui generis* properties of objects, i.e. properties not reducible to properties described in non-chromatic terms, e.g. in terms of physical attributes of objects' surfaces or reflected light, etc. A primitivist about colour – in particular, one who is also a realist about colour – will admit that the nature of perceived colour does not correspond well to the microphysical properties of objects' coloured surfaces or to the properties of coloured light, and so cannot be identified with those properties. However, he will insist that colours are nonetheless 'out there', inhering in worldly objects, in just the way that the natural view requires.

Confronted with the correlation between perceived colour and the structure and functioning of the visual system, the primitivist might claim that the visual system has evolved precisely so that it responds to the real colour properties of objects. In other words, the nature of colours has determined the nature of the visual system, and not *vice*

versa. However, one problem for the primitivist is that light mediates colour experience in some sense – certainly, in physical terms it mediates between the ostensibly coloured object and the subject's visual system. As such, if primitivism is to appeal to a (causal) determination relation between real colours and the structure of the visual system it must show how a physical system can be determined by something – colour – that is not itself reducible to the physical. If, on the other hand, the primitivist ignores the correlation between perceived colour and the structure of the visual system, his theory looks at best unexplanatory and at worst perverse.

The problem is compounded by the observation that different organisms (and indeed colour-blind human subjects) have visual systems with different structures. These differences are evident even at the retinal level, never mind any downstream processing: some colour-blind subjects have two kinds of colour photoreceptors (cones), while pigeons and goldfish have four (Thompson 1995: 152). By analogy with normal (trichromatic) human vision, it has been suggested that creatures with a greater variety of colour photoreceptors will be aware of a correspondingly greater number of colours.⁷⁶ Crucially, it has been suggested that these 'extra' colours would not amount simply to finer distinctions within normal human colour space but would partly define a wholly different colour space with colours quite different from those accessible to normal humans.⁷⁷ These 'spaces' might indeed have a different number of dimensions: whereas the human colour space with its three dimensions of green-red, blue-yellow and black-white (brightness) is readily visualised by analogy to the three dimensions of space, the colour space for a tetrachromat like a pigeon might be defined along four-dimensions. This would render the respective colour spaces incommensurable insofar as it would not be possible to obtain a one-to-one mapping of colours (or colour differences) between the spaces (Thompson 1995: 145-152).

It seems, then, that primitivism can accommodate these claims only by admitting that the colours of worldly objects might be different for different for different observers; that

⁷⁶ Some animals have an even greater number of receptor types: an extreme example is the mantis shrimp, some species of which have 12 different photoreceptor types, although their apparent failure to make commensurately fine colour discriminations is attributed to likely differences in post-receptoral processing, and is taken to undermine the thought that they perceive a proportionately greater variety of basic colours than creatures with fewer receptor types (see Thoen et al. 2014).

⁷⁷ Note that, crucially, the different visual systems here might be sensitive to the same range of visual stimuli, so that differences in visual systems and/or colour experiences cannot be attributed merely to differences in what organisms are visually sensitive to.

what a human sees as red a pigeon might see as something different and possibly unfamiliar to humans. This seems antithetical to the very purpose of primitivism as it is usually presented, namely as a way to defend from the failure of attempts at physicalist reduction the claim that colours are out there inhering in worldly objects.

In this point about inter-species differences in visual processing we might find a bridge between claims for subject relativity ^{general} and subject relativity ^{particular}. I have argued that subject relativity ^{general} is best supported by empirical evidence for what Pautz calls bad external correlation and, conversely, good internal correlation. The appeal to interspecies differences now assumes good internal correlation and uses this to challenge primitivism with the resulting conclusion that the same objects can simultaneously possess different colours for different organisms. As we saw earlier, there is good independent evidence for marked intraspecific but intersubjective variation in hue perception, and primitivism would seem similarly ill-placed to account for this. Nor is it clear how primitivism can accommodate the colour contrast/adaptation phenomena which illustrate intrasubjective perceptual variation. Primitivism would therefore seem unable to explain either subject relativity ^{general} or subject relativity ^{particular}.

6.5 Subject-relativity and the simple metaphysical picture

How are we to think about the nature of colour given the conjunction of the natural view and subject relativity, and the explanation offered in terms of the simple metaphysical picture? Earlier I suggested that it presents us with two key questions:

- 1. How can we make sense of qualities' being out there in the environment and yet also determined in part by the subject?
- 2. What does this entail for other qualities of objects that do not seem to be subject-relative, and indeed for the objects themselves?

I have already noted that one broad theory of colour is aimed precisely at answering the first question. This is relationalism about colour, which holds that colours are "constituted in terms of some relation between (inter alia) objects and perceivers" (Cohen 2004: 452). The most widespread form of relationalism is *dispositionalism*, whereby an object's being red is identified with its disposition to *look* red. The relevant disposition is typically explicated in causal terms, so that a disposition to look red is explained as a disposition to cause experiences as of redness. Formulated in this way, the dispositional account clearly conflicts with my defence of the natural view via the simple metaphysical picture, since the latter denies that experiences are caused by their objects; indeed, it denies that there *are* experiences in any substantial sense.

A non-causal account of an object's being disposed to look red looks no more plausible: if it is not a disposition to cause an experience as of redness then it is hard to see what other options remain (surely not a disposition to change in such a way as to *become* red). Intuitively, when we see an object's colour, we do not see it as doing anything – rather we see it simply being coloured. According to the natural view, the way an object looks is just a matter of how things are with the object and its surroundings (the conditions of lighting, its position within the perspective of the subject, etc.).⁷⁸ From the discussion of physical illusions we can appreciate that there might be various ways for an object to look red – for us to see it in red light or through red-tinted spectacles, for example – but the paradigm kind of case would be one in which the viewing conditions are favourable and the object *is* red. To look red in such a case is for the object seen to *be* red.

A related worry about the application of the dispositional account in the context of the natural view is that the dispositional account of colour arguably fails to locate colour 'out there' in the right way. As noted earlier, it is implicit in the natural view that the various perceptible qualities of worldly objects are 'on a par' in perceptual experience: it is the object itself and its various qualities that simply feature as the constituents of phenomenology. The tomato's redness is seen in the same way as its roundness – just by being there, qualifying the tomato itself. Equating colours with dispositions, meanwhile, potentially opens up the primary/secondary quality or objective/subjective distinction alluded to in the second question above. Certain 'primary' qualities of objects are accessible to vision, shape being the obvious example. To see an object's shape is presumably not to see an object's disposition to look round, say; it is just to see its roundness. We could assert this consistently with diverse theories of perception, including

⁷⁸ We should note also that 'looks' is potentially ambiguous between phenomenal, comparative and epistemic readings so that, for example, the reflective surface of a wooden table in bright sunlight might look (phenomenally) almost white, but simultaneously look (epistemically) like it is brown (see Dretske 2000). In the current context, 'looks' is perhaps most appropriately taken to have its phenomenal sense.
those that postulate sense-data or qualia, but it is clearly essential to the natural view. After all, according to the natural view, it is the object itself and its qualities that constitute the experience's phenomenology.

Thus, one reason to reject the dispositional view of colour is that it would seem to leave colours 'out there' in a rather different sense from the sense in which so-called primary qualities are out there. The latter are paradigmatic examples of 'objective' qualities of worldly object – qualities specifiable using methods available to any subject and without essential reference to how they appear perceptually to subjects. According to the natural view it is just these qualities themselves that contribute constitutively to the phenomenology of experience. Colours (and other 'secondary' qualities like taste and sound) are likewise supposed to contribute constitutively to phenomenology in the same way, as qualities inhering in the objects themselves. Dispositional accounts, however, leave colour rather more on the 'subject side' than this suggests. This may force them to admit two different uses of 'colour': on the one hand to denote the object's disposition (or the categorical property underlying this disposition) and, on the other hand, to denote the phenomenal quality that the object is disposed to produce. Trying to collapse the two into a single sense might seem doomed to failure.

If dispositionalism seems unsuited to underpin the natural view, subject relativity nonetheless seems to entail some sort of relationalism about colour. After all, it does indeed appear that an object's colour is constituted by the relations between that object and the perceiving subject. But if dispositionalism is not the way to flesh this out, how should we go about it? Is there a way to reconcile subject relativity, and so relationalism, with the natural view's insistence that colours, etc., are straightforwardly 'out there'?

I will argue that the answer is 'yes', but that doing so will extract a high price in terms of our commitment to realism and certain customary assumptions about the metaphysical structure of the world. It will be obvious that an affirmative answer will necessitate some sort of claim for the 'mind-dependence' of worldly objects or qualities. Nonetheless, I will show that we can recast such 'mind-dependence' in terms of a broadly physicalist ontology that eschews talk of 'minds' altogether. The ontology will prove simple – indeed, monistic – and although it is controversial, it has some independent support both philosophically and empirically.

This metaphysical account will occupy the following chapter. For now, in concluding this discussion of perceptual variation, I will sketch out how we get there by conjoining the natural view and subject relativity. Again, the key bridge between the two theses is the simple metaphysical picture of experience. Recall that what it is for an object or quality to be a constituent of an experience is explicated via the simple metaphysical picture in terms of its occupying a certain place in the causal process that is the object's being perceived by the subject. Now let us recast the claim that colours are subject-relative as the claim that they are experience-dependent, i.e. their nature is determined at least in part by the nature of the perceptual experience with the process, the experience-dependence of colours becomes the claim that colours are determined by the perceptual processes in which they participate (as qualities inhering in objects).

This of course inverts our customary understanding of the relationship between objects and the processes in which they participate. We would normally think of processes as being constituted by – built out of, to put it crudely – the various objects that participate in them. Constitution here is supposed to be allied to a form of determination, so that a given object has its nature determined by the antecedent nature of its constituents. There may be other forms of determination at play, such as causal determination, but the relevant determination relation here is precisely the constitutive one. Central to the claim about constitutive determination is just this notion of metaphysical priority: given any object divisible into parts, the nature of those parts is metaphysically prior to, and thereby determines, the nature of the whole object.

We might call this customary metaphysical assumption *bottom-up determination*. By contrast, I am suggesting that applying the simple metaphysical picture of experience to cases of subject relativity or experience-dependence points to a converse picture we might call *top-down determination*. The claim, then, is that objects' colours are subject-relative or experience-dependent because their nature is dependent on, and determined by, the nature of the larger (perceptual) process of which the objects are constituents. We could apply this theory to each of the three phenomena described above, namely simultaneous colour contrast, intersubjective variation in hue perception, and the structure of colour

space. In each case it is the nature of the whole perceptual process or relation between subject and object that determines the nature of the colour perceived.

Certainly, this view abandons the realism claim, i.e. that worldly objects exist and have their perceptible qualities independently of their being perceived. Consequently, it also abandons the explanatory virtue of naïve realism, by which the experienceindependent qualities of the object determine the phenomenology of experience. On the view I have presented the explanation goes at least partly in reverse; the properties of the scene are to be explained at least partly by reference to the whole perceptual process of which the scene is part. At this point, the view might also seem to risk incoherence. The simple metaphysical picture, after all, purports to explain experience in terms of a process that, implicitly, is defined over its (physical) constituents. Now, instead, it is claimed that the nature of the constituents is at least partly to be explained in terms of the process of which they are parts. So we seem to be seeking simultaneously to explain a process (experience) in terms of its constituents, and those constituents in terms of that same process. This, however, is an inevitable consequence of the notion of top-down determination defended here.

Some will further object that the view divests itself of its naivety as well as its realism, on the grounds that it is part of pre-philosophical common-sense that objects' qualities are constitutively independent of our perceptions. I think there is a common-sense insight in the vicinity here, but that it has more to do with a grasp of our causal powers and their limits. So striking something affects things but merely looking at them does not. What is in at stake in the case of colour perception is not, however, a matter of our causal relations to worldly things and properties but our *constitutive* relations. And here I think there are no pre-philosophical intuitions to be had. The 'naïvety' to which the view is faithful is the sense that, in perceiving, we are simply presented with the world as it is (whether or not 'how it is' needs to be supplemented with 'for us').

I should, however, note the limitations of the claim supported by the evidence of the colour phenomena described so far. Taken by themselves, these phenomena entail only a limited rejection of the realism claim: they show at most that *some* qualities (specifically, the colours) of worldly objects are constituted partly by the perceptual processes in which they participate. They do not tell us whether the same should be said of other qualities of

perceived objects, although they do suggest by analogy how we might look for the subjectrelativity of other qualities. I will leave it to the next, and final, chapter to consider the fuller metaphysical implications of subject-relativity for our view of objects' primary qualities and for the objects themselves. Pressure to expand the scope of the top-down determination relation to objects and their primary qualities arises from the implicit requirement of the natural view that the various qualities of perceived objects, both 'primary' and 'secondary', have a certain parity as qualities genuinely out there, inhering in worldly objects. Cashing out what this parity amounts to will prove challenging, but it will ultimately require that worldly objects and their primary qualities too are experiencedependent insofar as they do not have their nature independently of the wider processes in which they participate.

Note also that although I have rejected the realism claim, the realism ruled out is of a restricted kind, framed in terms of the constitutive independence of perceived properties from their being perceived (their participating in the perceptual process). Set against this, the simple metaphysical picture presumes an ontology of physical processes of which objects with their properties are constituents. Indeed, it needn't presume anything else – no especially 'mental' entities or properties except those that are defined over physical processes. As such, whereas the conjunction of the natural view and subject relativity (or experience-dependence) might seem to point to a sort of idealism in which objects and qualities are dependent on minds for their existence, the upshot will instead be a radical holism about the physical; one that results from reversing the normal assumption of ontological priority from parts to wholes. Neither idealism nor a notion of realism framed around 'mind-independence' is straightforwardly applicable to the view sketched here. I will say more about this in the next chapter.

CHAPTER 7 - METAPHYSICAL IMPLICATIONS OF THE NATURAL VIEW

7.1 Introduction

In the first part of this thesis I have sought to defend the view that perception acquaints us directly with the world around us. Key to that defence has been a *simple metaphysical picture* of perceptual experience, as nothing more or less than the state of affairs that is the object's being perceived by the subject, or the subject's perceiving the object. I have further suggested that this state of affairs can be described in physical terms as a causal process involving, *inter alia*, the perceived object, the subject and relevant perceptual intermediaries. I showed how the simple metaphysical picture can not only account for the sort of direct perceptual acquaintance with our environment that is expressed in the natural view, but can also explain other forms of sensory experience or phenomenal consciousness, perceptual or otherwise. As such, the simple metaphysical picture possesses an explanatory scope beyond the straightforwardly perceptual cases that inspired it.

In this chapter I will seek to broaden the scope further and consider what acceptance of the simple metaphysical picture might mean for our view of the world more generally, beyond a narrow concern with sensory experience. This should be of general interest, but its importance to the current thesis stems from two main considerations. First, the simple metaphysical picture, although proposed to account for sensory experience, is not about some special, restricted class of 'sense experiences', construed as peculiarly 'mental' states or events. Indeed, it is a defining feature of the view that there *are* no such states or events. The simple metaphysical picture precisely collapses claims about experiences into claims about the relations among subjects and their environment, all construed broadly in physical terms. That is, it claims that sensory experiences can be understood as physical processes which are distinctive purely on account of their relating subjects and other worldly objects in certain sorts of ways (involving light, eyes, etc.). Other more or less familiar physical processes may not be perceptual or experiential, but that is not because they lack some metaphysically special ingredient, but simply because they are not the right sort of processes relating subjects and objects. Thus, if perception leads us towards any more or less fundamental metaphysical claims about the relationship between processes and their parts, we might expect those claims to bear directly on all manner of processes and not just the specifically perceptual ones.

That, then, is one reason to tease out the wider implications of the simple metaphysical picture. The second reason is more specific, and has to do with the phenomenon of subject-relativity or experience-dependence described in chapter 6 above. It is here, if anywhere, that a feature of sensory experience seems to drive our metaphysics into unfamiliar territory. Taken at face value, it seems indeed to demand a fundamental reversal of our customary assumptions about the relationship between ourselves as subjects and the world that we encounter with our senses.

I brought out experience-dependence using colour as perhaps the most striking example. Colours' experience-dependence was shown to manifest itself in broadly two ways. Firstly, the very nature of the colours, including their mutual relations of similarity, cannot be explained solely by the nature of the coloured objects out in the world, but only by reference to some physiological properties of subjects. And, second, perceived colour is not only dependent in this way on subjects in general, but also varies among different individual subjects.

These findings force us to reject naïve realism as I have formulated it, in which case we must abandon at least one of its constituent theses, i.e. the natural view or the realism claim. An obvious choice would be to reject the natural view and hold on to the realism claim, and this is the kind of move that has historically been favoured by philosophers when confronted with seeming problem cases for naïve realism. This choice is by no means forced on us however – we can choose instead to retain the natural view, in which case we must reject instead the realism claim. In other words, we might embrace the conclusion that perceptible worldly objects are, at least in some respects, dependent on their being perceived, or on the subjects that perceive them.

This is likely to spark fears of incipient idealism. However, my aim in this chapter is to show that we can in fact reconcile the subject-relativity or experience-dependence of worldly objects with a broadly physical ontology. To that end, I will argue that the simple metaphysical picture allows us to understand the subject-relativity of worldly objects or qualities in a way that preserves our sparse, broadly physical ontology and avoids some of the more unappealing idealist consequences. In order to do so, however, our metaphysical picture will inevitably become less simple, and will incorporate some claims that are less than intuitive.

Key amongst these metaphysical claims are the following:

- 1. A fundamental metaphysical structuring principle of *top-down determination* in which the nature of individual objects or events is determined by the nature of the whole processes in which they participate.
- 2. A *process monism* in which the whole universe from beginning to present is fundamental.

The sceptical reader might be tempted to infer from these metaphysical conclusions that the thesis amounts to a *reductio ad absurdum* against the natural view. A *reductio* is not my intent of course. This chapter will certainly aim to show what bullets we may be forced to bite if we are to defend the natural view of perception via the simple metaphysical picture. It is important to note, however, that some or all of these are bullets we may already have to bite as a consequence of our best physical theories.⁷⁹ We therefore have reasons to accept some or all of our major conclusions besides the reasons provided straightforwardly by perception. Whether these other reasons are wholly independent of the perceptual considerations, and how forceful they are in their own right, are questions beyond the scope of this thesis, although one might expect that two sources of evidence for the same or similar metaphysical picture share some underlying relationship.

I should emphasise, then, that my aim is not to offer a detailed defence of the wider metaphysical picture proposed. Some of its constituent claims – monism in particular – have been the subject of considerable debate. However, those debates have been conducted largely without concern for the sorts of perceptual considerations that motivate the natural view and simple metaphysical picture. Whether monism or top-down determination are plausible or defensible independently of these perceptual considerations is well beyond the scope of this thesis. On the plus side, we might view the claims in this chapter as offering a novel argument – an extended inference to the best explanation – in

⁷⁹ Most obviously, the majority of interpretations of quantum mechanics support a holistic view of physical reality, in which the states of what look like independent objects are actually determined by the wider states or processes in which those objects participate (see Esfeld 1999, Schaffer 2010: 51-5).

favour of monism, or specifically the form of process monism described. Given, however, that it rests ultimately on a controversial account of perception, this might give limited succour to the monist approaching from a different direction.

The metaphysical project that this chapter seeks to propose, if not complete, might be viewed as an attempt to answer the following question: *What must the world be like for the natural view to be true*? Posing the question this way presents the natural view as a source of problems (problems arising from illusion, hallucination, cognition, etc.) whose solution requires some crafty metaphysics. In other words, it makes it appear as though we must contrive our metaphysical picture of the world in order to fit the claim that perception acquaints us directly with the world. While this is not inaccurate as a summary of the philosophical task, there is a more positive way to present it. From this more positive perspective the guiding question becomes this: *Assuming that the natural view is true, what can we learn from perception about the nature of the world*? This reframing of the question acknowledges the central virtue of the natural view that perception as a source of *data* for our metaphysical picture rather than a source of *problems* to be fixed by our metaphysics. In other words, the natural view licences us to take our perceptual experiences far more at face value, as revealing the true nature of the world.

In this chapter I will therefore describe what appear to be the key aspects of the world's nature as it is revealed to us in perception. Given the various puzzling perceptual phenomena considered in the first part of this thesis, there would seem to be two such aspects. Firstly, as already described, the world is (in at least some respects) subject-relative. Second, we can be acquainted with the past in much the same way as we are perceptually acquainted with the present. This second aspect therefore gives us cause to consider what perception reveals about the temporal structure of the world, given my earlier claims that such perceptual phenomena as perceptual memory, hallucination and cognitive penetration involve direct awareness of objects previously perceived.

The two main sections that follow will consider these aspects in turn, with subsections that tease out the implications of these aspects for our metaphysical picture of the world and our perceptual relation to it.

7.2 Worldly objects are subject-relative

Clearly, the most distinctive and surprising aspect of the world as revealed by perception is its subject-relativity. Taking our experiences at face value we discover that colours, for example, are out there inhering in objects and yet are subject-relative or experience-dependent, in the ways described in the previous chapter. A central challenge in this section will therefore be to explore further how this subject-relativity or experiencedependence ought to be understood if we are to explain the nature of perception consistently with the natural view and the simple metaphysical picture. In so doing I will seek to outline a broader metaphysical picture adequate to accommodate both the natural view and the subject-relativity of perceptible qualities such as colour.

Crucially, I will suggest that the simple metaphysical picture offers a way to reframe subject-relativity in a way that spares us from idealism and opens up an interesting line of metaphysical inquiry into the nature of the physical. In so doing, it holds the promise of a metaphysical picture that is not only adequate to explain how things appear perceptually but also possesses virtues we would aspire to find in any credible theory. These will include not only empirical adequacy but also broad explanatory scope and simplicity/parsimony.

7.2.1 From experience-dependence to top-down physical determination

As I described in chapter 6, the subject-relativity of worldly objects manifests most obviously in respect of their colour. We found that the subject-relativity of colour takes broadly three forms. Firstly, there is *intra*subjective variation – variation in the apparent colour of an object when seen by the same subject but under different viewing conditions (the Rubik's cube illusion described in chapter 3.5 illustrates this). Second, there is *inter*subjective variation in the apparent colour of objects perceived under the same conditions (variation in unique hue perception was cited here). The third aspect of colour's subject relativity is apparent in the nature of colour in general (as perceived by normal human subjects, notwithstanding the intersubjective variability already noted), whereby by the structure of colour 'space' – the pattern of resemblances among colours – proves to be better correlated with neural processing internal to the subject than with physically describable features of the ostensibly coloured object. The first two forms of subject relativity are therefore manifest in differences between particular colour experiences (exhibiting what I called subject relativity ^{particular}), while the third is apparent from the nature of colour in general (exhibiting what I called subject relativity ^{general}).

I argued that subject relativity of either sort need not be taken to imply that it is something in the nature of the subject alone that determines the nature of the object or its colour qualities; crudely, that something in the subject 'reaches out' and fixes how the world is colour-wise. Rather, we might usefully think of it as experience- or awareness-dependence rather than subject-dependence, i.e. it is the nature of the *experience* and not merely the subject that determines the nature of the object or its colour qualities. Such a suggestion remains apt to provoke accusations of idealism, as if worldly objects and/or their qualities were to be understood as mind-dependent. However, in my exposition of the natural view I have eschewed talk of 'minds' or 'mental states' and explicitly rejected the notion of experiences as inner mental states or episodes, taking them instead to be nothing more than the states of affairs in which the subject and worldly objects are linked perceptually – the object's being perceived by the subject, or the subject's perceiving the object.

Left in these terms, the account might still seem to present perceptual awareness as a mysterious, 'world-constituting' relation. However, in setting out what I have called the simple metaphysical picture, I said that the notion of perceptual experience as nothing more or less than the perceptual relation between object and subject admits of a redescription in physical terms. To redescribe the perceptual relation in this way is just to say what is going in perception and to say it in terms of the very sorts of things we can encounter perceptually.⁸⁰ The upshot is that we describe a process in which, typically and amongst other things, the object affects the subject appropriately via her sense organs.⁸¹

⁸⁰ Of course the physical redescription will include not just the things we encounter in the perception being described but will also make reference to things that we typically don't see *when* we see (e.g. our neural states) and properties of things that are typically not apparent to us when perceive (wavelengths of light for example). Nonetheless, these are all things that we know about ultimately through perception, whether or not our perceptual access to them is mediated by scientific apparatus. And they all seem to share a similar ontological status as worldly things in some ordinary sense. The contrast is with such things as 'sense data', 'qualia', 'experiences' etc., none of which seem to occupy the world in the same way, or indeed at all.

⁸¹ Typically, because we can imagine contrived scenarios in which an object affects the subject's brain in a way that bypasses the normal sense organs, but in so doing nonetheless elicits a kind of awareness of that object. Amongst other things, because the relevant experience might be one whose phenomenology is fully explicable only by reference additionally to causal links between the subject and objects other than

This physical description of the perceptual relation opens up a way of thinking about experience-dependence that explicitly steers clear of idealist implications or any mystery about the awareness relation. Specifically, it allows us to reframe experience-dependence as a dependence relation holding between part and whole – specifically, between the object and the wider perceptual process of which it is a constituent. We could say that the nature of the object, or one or more of its properties, is dependent on the nature of the larger process or event in which it participates. Conversely, the nature of the process determines, at least in part, the nature of the object.⁸²

This is, of course, quite the converse of our customary assumption about the determination relation between processes and their constituents, or parts and wholes more generally, namely the assumption that physical wholes and processes are 'built out of' their constituent parts and have their nature determined entirely by the nature of those parts (and perhaps some laws of nature governing the way those parts interact). This customary view amounts to I will call an *atomistic* picture of the world, one in which the fundamental metaphysical structuring principle is a kind of *bottom-up determination*, i.e. the antecedent nature of the parts or atoms determines the nature of the wholes and processes that those parts combine to form.⁸³

Crucially, we can find the same structuring principle implicit in what I called the explanatory virtue of naïve realism; namely, the theory's promise to explain what it's like to perceive – the experiential state of affairs that is an object's appearing some way to a

the one immediately perceived; I am thinking here of the kinds of perceptual learning and cognitive penetration described in chapter 5, in which the character of a subject's experience is determined in part by remembered objects as well as by objects currently perceived.

⁸² Determination is a stronger relation than dependence: for B to depend on A does not require that A necessitate B, whereas the latter is entailed is entailed if A determines B (Grimes 1991: 83). Determination is what I have in mind here, and a correspondingly strengthened notion of dependence is therefore implied in relevant contexts.

⁸³ I will make use of a fairly minimal notion of atomism as a foil for the alternative view that I will develop in this chapter. To be an atomist on this conception it is sufficient to hold that, given the set of all worldly things that are discernibly composed of (smaller) parts, one should look among those parts to find the entities that are metaphysically prior or basic. It is in this sense that metaphysical determination goes from the bottom up. According to a formulation by Schaffer, atomism is the conjunction of three claims: (1) the universe as a whole is not basic (i.e. the denial of monism); (2) there is a plurality of basic entities; and (3) the basic entities do not have parts (Schaffer 2010: 43-4). The first two claims together yield what Schaffer calls *pluralism*, while the third claim yields atomism as a more specific form of pluralism (ibid.). It is atomism rather than pluralism that I will take as the foil for my form of monism, since it is more explicitly consistent with the claim that bottom-up determination is the fundamental metaphysical structuring principle. One might otherwise be a pluralist and suppose that some of the metaphysically basic entities have parts whose nature is determined from the top down.

subject – by reference to the antecedent, experience-independent nature of the objects perceived. I have argued that certain features of colour perception, when coupled with the natural view, show that the explanatory virtue does not obtain at least in relevant cases. The natural view insists that colours are 'out there', inhering in worldly objects, and that the character of our experience is simply the character of the object (as seen under certain ambient conditions and from a certain perspective). Now, insofar as different ostensibly normal subjects perceive the same thing as having different colours, even under the same viewing conditions, it seems that how the object is (for a given subject) is a function of how it looks (for that same subject); effectively a reversal of the explanatory virtue, and so yielding what might seem to be a vice of the natural view.

Recast in terms of physical processes, as per the simple metaphysical picture, the same observation yields the view that, at least in some respects (obviously, colour), the nature of objects is determined by the nature of the larger processes (here, perceptual) in which they participate. At least in such cases atomism and its bottom-up determination claim do not seem to hold. We appear instead to be committed to what we might call a *relativist* metaphysical picture, where by 'relativism' in this context I mean to signal the claim that, in at least some respects, objects have their nature only relative to the wider processes in which they participate. ⁸⁴ In other words, the relevant aspects of an object's nature are grounded in, or obtain in virtue of, the wider state of affairs in which that object is related to some other objects and their properties.⁸⁵ Where the atomist view involves a structuring principle I have called bottom-up determination, the relativist metaphysics involves a structuring principle that we might call *top-down determination*.

It is important to note that what is being proposed is not a relation of *causal* grounding or *causal* determination holding between whole processes and their constituent parts. The determination in question is claimed to be at work wherever we find objects standing in certain causal relations to one another (i.e. as constituents of some process), but this is not to say that the process or physical state of affairs *causes* its constituents to be some way or

⁸⁴ I use 'relativism' to avoid confusion with what I have earlier called 'relationalism', i.e. the claim that perceptual experience is a relation. No connection with physical theories of relativity is implied.

⁸⁵ In using the notion of grounding to express some uses of the 'in virtue of' relation, and vice versa, I follow widespread current practice, as defended by e.g. Audi (2012). Whether 'grounding' itself univocally denotes a distinctive metaphysical relation is disputed (see e.g. Daly 2012 for some sceptical arguments) but I will put it to use in explicating a specific metaphysical thesis, such that different uses of 'grounding' or 'in virtue of' needn't be thought of as denoting the same fundamental metaphysical relation.

other. The grounding or determination relation is perhaps better thought of as constitutive rather than causal, although describing it this way is likely to lead to confusion, since we more customarily think of processes being constituted by their parts (causes and effects) rather than *vice versa*. If we were to talk of the top-down determination relation as a form of constitutive (as opposed to causal) determination we would imply that the parts are constituted by the whole rather than *vice versa*. That is not in itself an unreasonable way to frame the top-down determination claim, but it runs entirely counter to the ordinary talk of wholes as having constituents *qua* parts, a usage that I will preserve.⁸⁶ Thus, I will settle for talking of top-down determination, while stressing that this is not a form of causal determination.

7.2.2 Wholes, processes and objects

Before exploring the further implications of the posited top-down determination, it will be useful to say a little more to justify the appeal to *processes* as the determining entities, and indeed to clarify just what a process is understood to *be* on this view.

If we are wondering what a process *is*, a first thing to say is that the concern here is ontological rather than semantic – it has to do with what there is rather than with what our talk of processes (or, relatedly, causes) always means. To do the metaphysical work required by top-down determination, processes, in the relevant sense, must be genuine features of the world. This is not to say that all our 'process' talk picks out such genuine features, but conceptual analysis is not to the point here.⁸⁷ That is, in positing an ontologically robust notion of processes I don't seek to give an account adequate to the full scope of our causal concepts. Our concept of causation is itself highly debated and certainly outstrips the kind of process with which we are concerned – for example, we are

⁸⁶ One could perhaps defend the use of 'constitution' in both contexts on the grounds that material constitution is a symmetric relation, but this would lead to unnecessary confusion and is in any case arguably implausible (see Barker & Jago 2014: 190-1).

⁸⁷ See Dowe (2000: 1 ff.) for discussion of the distinction between conceptual analysis and what he calls 'empirical analysis', where the latter concerns what there is more than what our words mean or our concepts pick out. The distinction is unlikely to be a clean one, given that we use concepts to fix and express the domain of our interest. As Dowe admits, his own account strays from a narrow concern with a robustly physical notion of processes to an attempt to use that notion as the basis of a theory of causation (2009: 224). With the latter task he finds himself challenged by the familiar problem cases besetting any conceptual analysis of our ordinary concept of causation (see especially Dowe 2000: chs. 6 & 7). It is no part of my project to develop a theory of causation or even processes that measures up to our ordinary concepts.

apt to admit acts of omission and prevention as causes even when they do not contribute to the flows of energy that we think constitutive of processes (see, e.g., Dowe 2000: 126 ff.). For that reason, we should avoid seeking to *define* 'process' in terms of causation, even though we might think of processes as essentially causal. There is also a metaphysical reason to avoid this: if we treat processes as more fundamental than their constituents – as I am suggesting – then causes and effects, and the (causal) relations between them, turn out to be subordinate abstractions from the process itself.

For my purposes, the following picture will suffice to capture a relevant notion of process at a level of precision (or otherwise) adequate to the metaphysical project under consideration. The picture – which is hopefully reasonably intuitive – is of processes as more or less localised patterns of energy or matter evolving over time. More precisely, given a single region of spacetime, any non-disjoint distribution or pattern of matter or energy, of whatever form, located within that region will be considered a process. A process, on this view, is constituted by any continuous flow of energy occurring within a single, arbitrarily chosen region of spacetime.⁸⁸

Where we draw the boundaries on a given process will reflect our interests, although it will tend to reflect obvious discontinuities in form, as at the edges of those regions of relative stability we call 'objects'. Ultimately, such boundaries may be arbitrary, especially if, as cosmological theories involving a 'Big Bang' might hold, the whole universe from its origin to the present (and beyond) is a single process or evolving pattern of energy (this has monistic implications that I will consider in section 7.2.5 below). In an ordinary, everyday setting, we will single out those processes (including objects) that are most salient perceptually; in a scientific setting we may single out processes on the basis of some property or quantity (mass, say, or charge) that we wish to track over time through some changes in physical structure.

⁸⁸ The position I outline here is consistent with a 'conserved quantity' theory of causal processes, which appears to be the most defensible philosophical account of causation in terms of processes (e.g. Dowe 2009: 219-221; Dowe 1992; Salmon 1994). Conserved quantity theories offer a refinement of earlier theories that define causal processes in terms of transfers of energy-momentum (e.g. Fair 1979), amending the conserved quantities to those that, unlike energy or momentum, are conserved (energymass, energy-momentum) and incorporating other physical quantities, such as charge, that are also conserved in interaction irrespective of the relativistic frames of reference (Dowe 2009: 219). The resulting notion of process respects a key philosophical distinction drawn between genuine causal processes and 'unreal sequences' or 'pseudo processes', where the latter include such phenomena as shadows and moving spots of light (Salmon 1984: 141-143.).

I have suggested that processes, so understood, are metaphysically more basic or more fundamental than their constituent elements. This allows us to make sense of the experience-dependence of objects' colour properties, by way of identifying the perceptual experience itself with the perceptual process of which the coloured object is part, as per the simple metaphysical picture. Adding experience-dependence to the simple metaphysical picture generates an account in which the key metaphysical structuring principle is what I have called top-down determination, i.e. determination of the parts by the whole.

I was anxious in the introductory chapter to show that the link between the natural view and the simple metaphysical picture is itself a fairly natural one, even if it falls short of mutual entailment. We might therefore recall that the appeal to processes falls out of the natural view allied to some familiar empirical facts about perception: in short, we discover by looking at the world (with the aid of scientific instruments in some cases) that the objects we perceive are perceived only in virtue of the fact that we stand in certain forms of physical relationship with those objects; relationships that can be described as processes. When we investigate what is involved in perceiving, that is all we find – the objects, ourselves and certain intermediary entities. (That we don't necessarily *perceive* those intermediary entities while perceiving the object is, as I noted earlier, neither here nor there.) It is these observations that lead to the simple metaphysical picture whereby perceptual experience is understood as just the state of affairs or process that is the subject's perceiving the object's being perceived by the subject.

By this route, I hoped to have allayed the fear that an appeal to processes is unmotivated or *ad hoc*. However, a further worry might be felt to arise if we use the simple metaphysical picture to recast experience-dependence as a form of top-down determination. We might now wonder whether processes are merely one species of whole amongst a larger number to which top-down determination might apply. For instance, what about the more familiar whole/part relation holding between objects (at a time) and their spatial parts?

This worry is easily resolved if we recognise – as hinted above – that objects, *qua* concrete entities existing at a time, can themselves be considered as a special kind of process. While we might think typically of processes as dynamic and ever-changing,

concrete objects in a narrow sense – crudely speaking, lumps of matter – can be viewed as just those regions of processes that maintain a relative stability of form over time.⁸⁹ Thus, although we can treat objects as parts of processes, we can also treat them more or less in isolation as processes in their own right. Nonetheless, there appears to be no perfect discontinuity between objects and the wider processes in which they participate, or in which they appear; no discontinuity, that is, which would justify asserting some sharp ontological boundaries around them.⁹⁰

Allowing that objects, in the narrow sense, are themselves processes, or structurally stable parts of wider processes, offers a way to accommodate what is perhaps the more customary and intuitive way to frame questions of whole/part relations, i.e. in terms of objects and their parts. Understanding objects as processes also helps to avoid a worry that might otherwise attend the claim that wholes metaphysically determine their parts. If we were to treat any arbitrary mereological sum or fusion as a whole object we would be faced with claims of the following sort: that the colour of a leaf (for me) is determined by the perceptual process holding between the leaf and me, together with any arbitrarily chosen region of spacetime, such as a solar flare on Alpha Centauri 40 million years ago. That would seem to follow if top-down determination holds between any mereological sum and its parts. Of course, we might come up with some way of showing that the solar flare is merely a 'silent partner' in the determination relation. But the appeal to processes – single regions of spacetime in which there is a continuous flow of energy, in whatever form – removes any niggling arbitrariness, beyond the delineation of the spacetime regions themselves.

In fact, if we take the view that the universe is just one single process, any worries about disjointed mereological fusions and the like simply fall away. Ultimately, the 'top' in all top-down determination would thus turn out to be the universe as a whole, history and all. We could work our way to this conclusion from the bottom up, as it were. Thus, if we take any set of worldly objects linked as constituents of some localised process, the topdown determination claim demands that each of those objects is determined in some measure by all of the others. If we further assume that the localised process under

⁸⁹ Just which regions of a process are seen as possessing sufficient stability over time to merit being called 'objects' will of course depend on the spatial and temporal scale at which the process is viewed.

⁹⁰ Schaffer finds, in the evident arbitrariness of the 'joints' at which we carve nature, not to mention the vagueness of the joints themselves, intuitive support for the monism that I will discuss in section 7.2.5 below (Schaffer 2010: 48-9).

consideration is not causally isolated from the rest of the universe (as follows if the universe is itself a single process), then each of the objects in the localised set is dependent in some measure on every other object in the universe.⁹¹

The upshot would be what we might call a *process monism*, in which there is just one fundamental entity, namely the whole universe from its beginning to the present. I will indeed endorse process monism as the best way to accommodate experience-dependence to the natural view, and I will discuss the merits and implications of such a process monism in section 7.2.5 below.

However, much more needs to be said to justify such a move. We have reached process monism by way of the following key claims:

- The claim that in perceptual experience we are directly aware of our environment such that it is the environment itself that is the bearer of phenomenal character (the natural view).
- 2. A simple metaphysical picture that underpins the natural view, namely the claim that what we call perceptual experience is nothing more or less than the state of affairs that is the object's being perceived by the subject, or the subject's perceiving the object; a state of affairs describable in physical terms as a causal process involving, *inter alia*, the perceived object, the subject and relevant perceptual intermediaries.
- 3. The claim that we can account for the subject-relativity or experiencedependence of certain ('secondary') qualities consistently with the simple metaphysical picture by postulating a top-down determination relation holding between whole processes and their parts.

⁹¹ One might be tempted here to restrict the object's dependency to the set of objects that fall within its past light cone (and perhaps also its future light cone), since the past light cone circumscribes all possible historical causal influences on the object (and the future light cone its possible future causal influences). Such a restriction proves unenforceable, however. For one thing, it proceeds from the assumption that we can isolate one object as the 'focus' of various processes, an assumption at odds with the hypothesised ontological priority of the whole universe. Second, any constituent A of object B's past light cone might contain in its *future* light cone an object C which lies outwith B's past (or indeed future) light cone. If A is part of the whole process involving B then there is no principled reason to exclude C either.

4. The assumption that everything in the universe has a single causal origin, such that the universe, from its beginning to the present, is a single process.

We can arrive at process monism by conjoining the final assumption – that the universe is a single process – with the claim that top-down determination is the fundamental metaphysical structuring principle. What I have not done, however, is offer any compelling reasons to suppose that top-down determination *is* the fundamental metaphysical structuring principle, or even a wholly general one. As far as it goes, my discussion so far has shown only how we might invoke top-down determination to explain the (experience-dependent) nature of a narrow range of properties, including colours, an explanation that appears necessary only within the context of a narrow range of processes, namely perceptual ones. If the appeal to top-down determination is to have any plausibility, and is not to seem as a wholly *ad hoc* response to some perceptual puzzles, it will be important to show that it can be applied more generally, and without contradicting what we know about the world, e.g. by the scientific method.

What is missing from the account, then, are some arguments for the general (indeed, universal) scope of top-down determination, arguments that justify the move from experience-dependence to a more general top-down metaphysical priority, and perhaps ultimately to process monism. In the next section I will address this question of scope in more detail, and highlight the explanatory benefits of holding that top-down determination should be true of all properties and objects if it is true of any.

Even if it were possible to show that *everything* might be determined from the topdown, we would still stand in need of an explanation for why top-down determination nonetheless only *shows up* in respect of certain properties (e.g. colour) and under certain conditions (e.g. of perception). This is a further question I will consider in section 7.2.4 when I address the familiar distinction between 'primary' and 'secondary' qualities.

7.2.3 The scope of top-down determination: a dilemma

We might frame the scope worry as a dilemma. On the one hand, if top-down determination applies to only *some* objects and/or some of their properties and not others (for example, to objects' colour but not their shape) then we lack a unified metaphysical

account of objects' total nature and potentially undermine the central claim that an object's colour (etc.) is 'out there' inhering in the object in the same way that its shape (etc.) is 'out there'. On the other hand, if top-down determination applies quite generally to all objects and all of their properties we need an explanation of why this does not appear to be so; for example, why colour appears subject-relative or experience-dependent, but shape does not. That is, while the evident experience-dependence of colour motivates the top-down metaphysical picture, we must account for the countervailing phenomenological data whereby many properties of objects appear to be experience-*independent*.

My explicit concern has been the perceptual process and perceived objects, but the scope of the top-down determination claim is not only of general metaphysical interest but also important to the plausibility of the specifically perceptual application. There are two ways we might question the scope and generality of the claim that objects might be dependent on the processes of which they are constituents:

- 1. Does the claim concern only certain kinds of process?
- 2. Does the claim concern only certain kinds of objects or properties?

My purpose in introducing the claim is obviously to address questions concerning perception, so it is the perceptual process that is specifically at issue for this thesis. And I have motivated the claim solely by appeal to observations concerning the experience-dependence of certain perceptible qualities, namely colours. All the same, independently of my immediate concerns vis-à-vis perception the two questions above remain in play: does the claim concern *only* the processes that constitute perceptual awareness, and does it concern only colour? These more specific questions might appear independent. It could be that perception is the only kind of process to which the dependency/determination claim applies, but that the dependency concerns many different kinds of perceptible objects or properties; conversely, it could be that the claim concerns only colour. This opens up a range of possible answers to the scope question, not all of which are mutually exclusive.

Restricting the scope of the claim to only some processes and/or properties would be problematic for two reasons:

- The restriction sits uneasily with the claim, central to the 'naturalness' of the natural view, that colours (odours, etc.) are *out there* inhering in worldly objects.
- 2. The restriction, firstly to perceptual processes and, second, to a restricted class of qualities is metaphysically unexplained and therefore mysterious.

The two problems are closely related. To see how, consider *why* a scope restriction would sit uneasily with the claim that colours inhere in worldly objects. This is not because a scope restriction would involve *denying* that colours are out there – indeed that colours are out there is the datum that the dependency/determination claim is intended to accommodate. That colours are out there is therefore assumed, and any putative scope restriction on the dependency/determination claim might be posited to acknowledge that we have grounds for imposing the relevant metaphysical constraint only on some perceptible qualities and not others. The problem is really that the natural view takes colours to be out there inhering in worldly objects *in the same way* that less obviously experience-dependent qualities like size and shape are out there inhering in worldly objects. The worldly inherence of colour and shape are supposed intuitively to be 'on a par' in this respect, so that colours are out there if anything is.

The problem therefore is that the scope restriction on the top-down determination claim seems to mark a deep metaphysical difference between, say, colours and shape; one that is unexplained and mysterious. In other words, given that, *ex hypothesi*, colours *are* out there, the first problem enumerated above collapses into the second. Let us therefore look more closely at this second problem. We can usefully do this through the prism of the simple metaphysical picture of experience. Thus, if perceptual experience is nothing more or less than the obtaining of a certain relation between worldly entities, a relationship describable exhaustively in physical terms, then what is it about the perceptual relation/process in particular that imbues it with this peculiar property of determining the nature of some of its relata/constituents? Furthermore, why is this determining power exerted on only some qualities and not others?

Given the natural view, the crucial datum is this: whether we take objects to *bear* qualities or else to be *constituted* by them, the qualities all appear to be on a par inasmuch as they can inhabit the same regions of space and are accessible to us in the same way, i.e. via the senses. It would therefore be less mysterious if what would appear to be a fundamental metaphysically structuring principle – top-down determination – applied to all objects and qualities and not just a restricted class. In other words, we should hope, on grounds of explanatory simplicity and parsimony that top-down determination might prove to be a universal structuring principle applying to all objects and properties, or at least what we call *natural* properties, i.e. those 'ordinary' and broadly physical properties accessible to the senses and to scientific investigation.⁹²

Certainly, the universal scope of top-down determination is built into the process monism I described in the previous section (and which I will discuss further in section 7.2.5 below). On that account, there is really only one process, such that case if top-down determination applies to any process then it applies universally.⁹³

Nonetheless, claims for the universal scope of top-down determination, in whatever guise, will prove unpersuasive if we cannot explain why top-down determination reveals itself in respect of only *some* qualities and not others. It is all very well to give an explanation of the experience-dependence of colour in terms of top-down determination, but it will be of equal or greater importance to show how universal top-down determination can be squared with the appearance of experience-*independence* manifested by such qualities as shape and size. It is that task I will turn to in the following section.

7.2.4 'Primary' versus 'secondary' qualities

Why, then, might top-down determination only *show up* (in the form of experiencedependence) in respect of some properties and not others? I have focused on colours, but there are grounds for applying the top-down determination claim also to qualities accessed by different senses. For example, perceived olfactory qualities of taste and smell exhibit

⁹² There is some resemblance here to Lewis's notion of natural properties (1986: 60-1), albeit that on my view the relevant properties are clearly not to be thought of as intrinsic in the sense of their being possessed independently of how things are with the rest of the world.

⁹³ Of course, the universal scope of top-down determination would be just as consistent with a pluralism in which there were many processes, all of which were metaphysically fundamental.

similarity relations that, like those applying to colours, are not well explained by (correlate poorly with) the physico-chemical properties of the objects tasted or smelled, but correlate well with patterns of brain activity in the perceiving subject (Howard et al. 2009; Pautz 2011: 405-6). In addition, analogous to the intersubjective variation in colour perception, different subjects in some instances perceive the same chemical as having different odours, and this is explained by inherited variation in odour receptors (Keller et al. 2007; for a discussion of the seeming subject-relativity of smell see Batty 2010: 1153-4).

It is tempting to see these more or less obviously experience-dependent properties as comprising what have, following Locke, been called *secondary qualities*. For Locke these qualities, unlike the *primary* qualities like shape and size, are subjective inasmuch as they are to be identified with the powers of (e.g. ostensibly coloured) objects to produce ideas (of e.g. colour) in our minds (Locke 1993: 71-3). As such, secondary qualities are, for Locke, not really qualities inhering in objects, as they seem, but powers (ibid.). Primary qualities, on the other hand, are qualities that can properly be held to inhere in objects independently of their being perceived; crucially for Locke, primary qualities resemble the ideas they produce while secondary qualities do not (ibid.: 73).

Contrast the natural view – which holds all qualities to be straightforwardly 'out there' – with the Lockean account, in which all qualities are also out there, but in a different sense, namely as powers to produce certain ideas in the subject's mind. This leaves room for Locke to distinguish those (primary) qualities which resemble the resulting ideas from those (secondary) qualities which do not. However, given the natural view and my repudiation of 'experiences' *qua* distinctively 'mental' states, I have left no room for such a distinction – there is nowhere else for 'secondary' qualities to go if we deny that they are full-bloodedly out there.

It is in part this commitment to the co-location of 'primary' and 'secondary' qualities that drives the pursuit of a unified metaphysics. All the same, stipulating the universal scope of top-down determination purely in order to account for a limited range of perceptible qualities (the 'secondary' ones) is likely to seem a high price to pay, especially if we assume that other worldly properties (the 'primary' ones) are perfectly amenable to bottom-up, atomist explanation. It would therefore make a top-down metaphysics more plausible if we could offer one or both of the following: (a) a demonstration that what we think of as paradigmatic 'primary qualities' are in fact also subject to top-down determination; and (b) an explanation of why these 'primary qualities' lack the obvious subject-relativity afflicting colours and other 'secondary qualities'.

Satisfying the first desideratum seems likely to prove challenging at best, and I won't attempt it here.⁹⁴ In any case, satisfying the first desideratum will prove less crucial if we are able to show, as per the second desideratum, that we can explain the apparent subject-independence of the 'primary' qualities within the framework provided by a top-down metaphysics. I will shortly argue that we can do precisely this, so that the seeming primary/secondary quality distinction is demonstrably compatible with an all-encompassing top-down determination relation. Since we will thereby be accommodating the seeming experience-independence of 'primary' qualities, we will be able to make sense of our difficulty in satisfying the first desideratum – in showing that such qualities too are subject to top-down determination. The seemingly subject-independent nature of the 'primary' qualities thus becomes an observation to be explained rather than overturned.

The problem is, of course, not that 'secondary qualities' often appear different to different subjects (or to the same subject in different circumstances) while 'primary qualities' always appear the same. The appearance of 'primary qualities' also varies interand intrasubjectively. The difference is that variation in the appearance of some object's size or shape is typically attributable to differences in how things are out there in the environment, including the subject's spatial perspective. Thus, the round penny looks elliptical when viewed obliquely, and the tree looks smaller when viewed from afar. Unlike with certain colour phenomena, for example, we do not seem forced to explain differences in appearance by appeal to factors internal to the physical subject. So, while the nature of colours and odours might best be explained by reference to subjects' brain

⁹⁴ The relativist might take succour from our current best physical theories, which undermine the notion that objects have properties such as size, shape and mass absolutely. On the one hand, Einstein's Special Theory of Relativity famously demonstrates that properties such as those mentioned have no absolute magnitude but vary depending on the inertial reference frame in which they are measured. Objects will appear heavier or lighter, or larger or smaller, depending on whether, and how fast, they are moving relative to the measuring subject (e.g. Maudlin 2011: 45 ff.). On the other hand, quantum mechanics reveals that at very small scales (especially at sub-atomic scales) objects do not possess determinate values of such properties as mass, position or energy (bizarrely, determinate values may appear when the objects are subject to measurement, despite their demonstrably lacking determinate values *prior* to measurement) (e.g. Albert 1992: 73 ff.). While these theories in quite separate ways undermine the notion that objects have an absolute nature in respect of their 'primary' qualities, it is not obvious how, or indeed if, the underlying mechanisms bear directly on the top-down determination relation I am proposing.

states and might correspondingly vary from subject to subject, the same cannot be said of, say, shape or size. A corollary of this is that 'secondary' qualities appear essentially tied to perception in a way that 'primary' qualities do not. For example, colours seem essentially tied to (visual) perception: we have no way of getting at the nature of colour except through vision, a point reinforced by the failure, noted earlier, of attempts to identify colours with such physical properties as the wavelength distribution of light or the surface spectral reflectance of coloured objects.

What is needed, then, is an explanation of precisely why 'primary' qualities lack the obvious subject-relativity afflicting colours, and how this can be squared with the universality of top-down determination. If the latter is assumed then the question is why this determination *shows up* in respect of, say, colour and not shape. We can look at this from two angles: firstly, what is distinctive about secondary properties to make their experience-dependence obvious; and, second, what is distinctive about primary qualities such that they do *not* appear to be experience-dependent?

Take the first angle: why does the subject-relativity of secondary qualities show up? Continuing with colour as our example, we might note that it is a property that is essentially tied to vision – as Strawson puts it, "colours are visibilia or they are nothing" (2011: 142) – whereas shape, say, is accessible in various ways and via more than one sense modality. We can assess an object's shape using touch as well as vision, so that it does not seem tied essentially to how the object *looks* in the way that colour does.⁹⁵

But why should colour be essentially visual in this way? We can answer this precisely by appeal to top-down determination – that is, we can explain the object's colour nature by reference to the wider relation or process in which that colour appears. A distinctive feature of colour perception noted earlier is that our physiological response to objects' light-reflecting (/emitting/transmitting) properties exhibits a complexity or dimensionality exceeding that of the stimulus. If we view a range of light sources, each of fixed intensity but of one unique fixed wavelength, we are confronted with stimuli that vary along a single dimension. Nonetheless, our neurophysiological response in such a case is multidimensional in virtue of the structure of our visual system, including our possession of

⁹⁵ Recall from earlier discussion that colours resist identification with physical properties of objects' surfaces, such as their spectral reflectance profiles.

light-responsive cells in the retina which have different but overlapping sensitivity to different wavelengths, feeding into an opponent process in the brain's visual system. The upshot is that a stimulus varying along a single dimension (wavelength) appears perceptually in a way that can only be represented adequately in multiple dimensions.⁹⁶ From a first-person perspective, exposure to a light source of gradually increasing wavelength across the whole visible spectrum does not appear as mere quantitative change in a single quality, but as qualitatively discontinuous passage through the four unique hues and their various binary combinations. According to the natural view, the colour of the light itself passes through these qualitative 'boundaries' as its wavelength steadily increases or decreases.

As noted above, our perceptual sensitivity to other 'secondary' qualities may be similarly idiosyncratic, giving rise to obvious subject-relativity in those cases too (see Pautz 2014 for discussion). Smell and taste would be obvious examples, with genetic variation in olfactory and gustatory receptors taken to be responsible for intersubjective variation in smell and taste (Keller et al. 2007). Sounds are less obviously subject-relative; by contrast with the colour case, a one-dimensional change in wavelength (or frequency) corresponds to a one-dimensional change in perceived pitch (as wavelength decreases, pitch rises), although the correlation between wavelength and pitch is not perfectly linear, especially as one approaches the upper and lower limits of perceptible pitch, and can be influenced by changes in loudness (see O'Callaghan 2014: s.3.2.5).⁹⁷

If the idiosyncrasies of our physiological responses to stimuli are responsible for the more or less obvious subject-relativity of the relevant 'secondary' qualities, we might look for a corresponding *lack* of idiosyncrasy in our responses to 'primary' qualities like shape and size. To some extent this is likely to prove successful: we might reasonably anticipate

⁹⁶ Colour of constant lightness (i.e. just the hue and saturation components together) is often represented as a two-dimensional chromaticity diagram, but this diagram uses (i.e. *presents*) colours to show how different wavelengths of light appear to normal subjects; it does not *represent* colours. An adequate representation of spectral colour (as it appears to normal human perceivers) requires at least two dimensions to represent the range of similarity relations formed by the opponent pairs of blue-yellow and red-green, plus a third dimension to represent saturation. One could represent this, say, mathematically, without the *use* of colours. Note that the dimensions of hue, saturation and lightness are not wholly independent but interact in various ways (Thompson 1995: 50-1).

⁹⁷ Given that the subject-relativity of sound does not show up in everyday situations, it is interesting to reflect on why sounds have been so readily assimilated to colours as another kind of 'secondary' quality. Presumably the answer has to do with an intuition that sounds are essentially *heard* just as colours are essentially *seen*, coupled with a prejudice that anything heard or seen, *qua* 'sensations', must be in the mind rather than in the world.

a more or less linear and proportional relationship between changes in the actual size of an object and some magnitude of our physiological response to perceiving that object, all else being held fixed; certainly we have no reason to expect any radical discontinuities in our responsiveness to continuous changes in perceptible size. However, a complicating factor here might be the vulnerability of our judgments of size and shape to contextual and perspectival effects, a feature which is in fact crucial for our veridical perception even though it can also give rise to illusory experience and misjudgment (good examples might include optical illusions such as the Müller-Lyer described in chapter 3.5). As such, our responsiveness to shape and size is not quite as straightforward as all that. Nonetheless, we remain disinclined to say that size or shape are subject-relative, and the question remains why this is so.

We might appeal here to two further considerations. Firstly, there is the observation, noted above, that properties like shape and size appear accessible to more than one sensory modality – we can feel shape as well as see it, for example. Crucially, these senses do not seem to contradict one another, at least inasmuch as they exhibit the same dimensionality (felt changes of magnitude in X dimensions are matched by visible changes in X dimensions) and proportionality (felt increases in size are matched by relatively similar changes in visible size).

The second, and perhaps more illuminating, consideration has to do with the possibility of *standardisation* in our interactions with quantitative properties like shape and size. Objects can look of different size to different subjects or to the same subject under different conditions, but we can overcome these apparent differences through the shared use of some third thing – an object that serves as a standard by which we can measure the target object's disputed properties. What form this third object takes will depend on what property is to be measured – for size a measuring tape will often suffice. In this way, we can enforce or affirm our intersubjective agreement about the object's shape (or size, mass, etc.) through the use of intermediary objects that serve as measuring devices, so there are verifiable standards that guarantee consensus.

The standardising approach does not seem to work so well for 'secondary' qualities.⁹⁸ With 'primary' qualities like size and shape, the target of our standardisation – the object of our description relative to some mediating measuring instrument – still seems to be *that same quality*. We do not seem to have passed from talking about one quality – *shape as it looks*, perhaps – to talking about another – *shape as measured*. By contrast, when we standardise our description of an object's colour, say, we seem to find ourselves describing something else altogether: chemical properties of the object's surface, surface spectral reflectance, light wavelengths, etc.⁹⁹

We can help to make sense of this failure of 'secondary' qualities to submit neatly to measurement by viewing the perceptual relation from the top down, as it were. From that perspective, when we examine the perceptual relation we find it is composed of (or decomposes into) various participating constituents or sub-processes, including the physico-chemical properties of the object, the passage of light from object to subject, the neuro-physiological response, and so on. Of course, we also find the object's colour. Here is the crucial point: when we examine, on the one hand, the object's colour and, on the other, its physico-chemical properties, we are relating – physically, perceptually – to the object in quite different ways. In the first case we are simply looking at it; in the second we might also be looking at it, but now indirectly via various pieces of measuring apparatus (microscopes, spectrometers, etc.). On the top-down approach, it is in virtue of the first relation/process of direct visual awareness that the object presents as having a certain colour, while it is in virtue of the more complicated, apparatus-mediated relation/process that it presents as having its physico-chemical structure. The colour and the physico-chemical structure are not identical because they do not have their nature in virtue of all the same relation(s).

Even if this top-down, relational perspective accounts for the failure to identify or even correlate perceived colours with physical properties of their bearers (like physicochemical structure or surface spectral reflectance), the conclusion might seem to form an

⁹⁸ Note that we might agree standards for colour, say, in a slightly different sense: we might agree that the best way to guarantee that something which *looks* red is *really* red is to view it under optimal conditions, e.g. directly in broad daylight (see Allen 2010b). As such, for 'secondary' qualities as much as 'primary' qualities, we can draw a distinction between *how it appears* and *how it is*, albeit that the distinction has a pragmatic rather than absolute status.

⁹⁹ We might worry that colour therefore appears to 'float free' of these measurable properties of objects and light. I will return shortly to the question of whether colour is therefore compatible with a broadly physicalist ontology of the sort I have proposed.

inconsistent triad when placed next to two of my central claims. The natural view takes it as axiomatic that colours are properties that often inhere in objects – covering their surfaces or suffusing them. Meanwhile, I have explicitly committed myself to a broadly physicalist ontology, according to which everything belongs to the class of purely physical objects – the sorts of entities accessible to the senses and to the sciences and which, for example (and so far as we know), conform to mass-energy equivalence.

It is these two commitments – to the 'out-there-ness' of colour and to a broadly physicalist ontology – that might seem to be at odds with the claim that colours are not identical to physical properties of their bearers. I did say that the ontological picture suggested is only *broadly* physicalist to leave room for the existence of things and properties that might not themselves be considered to form part of any strictly scientific description of the world. In particular, I want to leave room for disagreement about whether colours should or could be part of any scientific description of coloured objects. The reason is precisely the failure to identify colours with any of the surface properties of the objects as described scientifically.

There is therefore a narrow sense of 'physical' in which we might deny that colours are physical properties of coloured objects. This is a different sense of 'physical' from the one I have used to describe my favoured ontology, which means only to rule out explicitly non-physical or immaterial objects and properties. In particular, I have denied the existence of distinctively 'mental' states or properties, unless 'mental' is taken merely to denote the sorts of (perfectly physical) relations among things that constitute awareness, thinking, and so on. The narrow sense of 'physical' has less to do with picking out a certain *type* of stuff than with marking off a certain range of methods for accessing objects and properties and describing them. These methods are what I have described as *measurements*, where these may involve the use of measuring tools as intermediaries between the objects and our senses. Colours are not physical properties in the narrow sense of 'physical' because they are not measurable in the same way as narrowly physical properties like spectral reflectance; colours are (directly) seen rather than measured. Of course, we could try to treat viewing subjects as measuring devices, ones which manifest certain brain states when confronted by appropriately coloured objects, or (assuming truthfulness and appropriate language use) make utterances like 'red' and 'turquoise'. However, looking at subjects 'from the outside' in this way is apt to mislead us about

colours: it might lead us to suggest that colours are the properties of objects that cause subjects to exhibit those brain states or behaviours. There is indeed a rather loose use of 'cause' in which we might say that the redness of the tomato (or my seeing it) caused me to utter the word 'red'. However, if a physicist were invited to describe the relevant chain of events the colour *per se* would likely not feature at all in the description. There would be reference to surface properties of the tomato all right, but these would be structural or chemical properties or relational physical properties (spectral reflectance). It is this outside perspective on the process that yields what we might call a narrow notion of the physical, i.e. precisely the intersubjectively invariant aspects of the observed system.

We could sum up the explanation of colour's experience-dependence or subjectrelativity in this way: it is in virtue of our being affected by objects' light-interactive properties in a certain way (or in virtue of the objects' affecting us a certain way) that those objects have the colours that they do. It is the idiosyncratic nature of our physiological response to light of different wavelengths that is responsible for much of the obvious subject-relativity of colour.

We must be careful here, however. The claim cannot be that our physiological response *determines* – or, worse still, *causes* – the object to be coloured. After all, the view proposed has it that the *whole experience* determines the object's colour. Although the claim is that this experience is just the whole state of affairs that is *my being affected by the object* or *the object's affecting me*, we need to avoid the impression that there is *first* the object, some physiological event in me, as well as some mediating events, and that these come together to determine or otherwise gives rise to the object's colour. This would be misleading for two reasons: firstly, the various discernible elements of the perceptual process are *already* together; and, second, the object's colour is already *there*, inhering in the object – the colour is already a feature of the perceptual state of affairs. So if we seek to explain an object's colour by reference to how that object affects us physiologically, it is not because the colour is in any way an *upshot* of this physiological effect, or that the object and its physiological effect on us are metaphysically prior to the object's being a certain colour.

It is rather the other way round: what is metaphysically prior is the whole experience or perceptual relation, in which the object appears with its colour. Here again we must

avoid making the converse mistake - that is, as much as we should avoid thinking of the physiological event in the subject causing the object's colour, we should also avoid thinking of the colour as causing the physiological event. The error in this line of thinking is perhaps less obvious. After all, there are three claims here that make such a thought quite compelling. First is the claim that the object causes the physiological event in the subject; second is the claim that the colour is out there, inhering in the subject; and third is the claim that the object's causing the physiological event is the process in virtue of which the subject sees the object's colour. The natural thought, given these three claims, is that it is the object's colour which causes the physiological event. That would an inference appropriate to a bottom-up metaphysical explanation of the experience's phenomenology, i.e. one compatible with naïve realism and its explanatory virtue of explaining the phenomenology by appeal to the antecedent nature of the object perceived. But it cannot be correct in light of the top-down determination claim and the observations about colour that led to it. These seem to show that the experience of colour is metaphysically prior to the colour itself, a claim cashed out via the simple metaphysical picture in terms of the metaphysical priority of the perceptual process over its parts.

But if an object's colour is supposed to be 'out there' just as much as its shape and physico-chemical properties (including those that we would certainly identify as causes of the subject's brain state), why can we not claim it also as a cause of the same?

This puzzle – like the converse mistake of treating colour as caused or determined by the physiological effect the coloured object has on the subject – arises if we are careless in conflating two different ways of reflecting on the perceptual experience: firstly, reflecting on it from the 'inside', i.e. the first-person perspective on how an object simply looks; and, second, reflecting on it from the 'outside', i.e. describing the experience as a physical process. These reflective perspectives on the experience are of course what yield the natural view and simple metaphysical picture respectively.

Conflating these two perspectives results in the following line of thought. The claim I am defending holds that perceived colours are really out there and that they are nonetheless determined by the whole perceptual relation. If we now (carelessly) recast that as the claim that colours are determined in part by their effects on the subject, then the same colours appear twice in the explanation of phenomenal qualities: first as the quality that

induces the physiological response, and again as the quality determined in part by that very response. In other words, the colour induces that which is partly responsible for its very nature. Behind the troubling implication of retrocausation, i.e. an effect determining one of its causes, is a blatant circularity.¹⁰⁰

The reason we should avoid conflating these 'internal' and 'external' perspectives on the perceptual relation is precisely that our relation to the object is different in the two cases: when we are simply viewing it ourselves as subjects, and when we are viewing it as a physical object affecting some other subject. In the second perspective, the object will still look (to us) coloured, but we now disregard that in favour of some other properties of the object – its surface spectral reflectance perhaps. As we look at it ourselves, the object is simply playing the role of a perceptibly coloured object; as we look at as the object of *someone else's* awareness, it is playing the role (for us) of a physical cause of that other subject's brain state. These are quite different relations between us and the object, and the object therefore manifests quite different properties in the context of each relation. This is precisely the point I made in the previous section, and it helps us to understand why the different properties of the object – its surface spectral reflectance, say, and its colour – so radically fail to correlate.

The puzzle we are left with, then, is not to explain how objects' colours *come about* given some antecedent facts about those objects and their effect on us. Rather, it is how to make sense of the relationships between the various things we find when we explore our experience. When we interrogate our experience – when, in other words, we look at the world – we find various objects, events, qualities and so on. Most of these are obviously and straightforwardly out there residing in our environment. Happily, there appear to be intelligible relations between the various things we find in the world (things we find by using our senses). In many cases these various things are measurable against some common standards, as described above for the 'primary' qualities. This allows us to maintain a certain 'objectivity', not just about those objects and their relevant qualities, but also about the (spatial, temporal) relations between them. In virtue of these common standards for spatiality and temporality (chronological time), we are able to describe the objects and their relations without essential reference to how they appear to us

¹⁰⁰ I consider the temporal aspect of this seeming circularity in section 7.3.2 below.

perceptually. We are, in other words, able to render those aspects of the world intelligible while keeping them 'at arm's length'.

We cannot, however, keep objects' colours (odours, etc.) at arm's length because of their distinctive dependence on the whole perceptual process in which our idiosyncratic physiological response is the upshot. There is a glaring mismatch between the response of our visual system to objects' surface properties (spectral reflectance) and the response of arm's-length measuring devices (spectrometers, say). Chiefly, the mismatch has to do with the peculiar dimensionality of our physiological response that we noted above, a dimensionality that seemingly results from the opponent processing occurring in our visual system. One might object here that if arm's-length measuring is all we are worried about we could surely construct some machine that recreates opponent processing of spectral light in the same way that the human visual cortex does. This would, however, be problematic, for two reasons. Firstly, calibrating a colour-measuring machine of this sort – setting it up to return values of pure green, red, etc. for given inputs – would be quite arbitrary, just as we cannot settle non-arbitrarily which human subjects perceive pure hues 'correctly'. Second, and relatedly, the output values of the machine would remain unaligned with common standards for measuring spatiality and temporality. As such, they would add nothing more to the intelligibility of colour than is provided by the colour talk of human subjects.

7.2.5 Process monism: prospects and problems

In section 7.2.3 I noted the explanatory advantages of taking top-down determination to have universal scope as the fundamental metaphysical structuring principle. I noted that we might defend this universal scope by appealing to the claim that the entire universe, from its beginning to the present, can plausibly be considered a single all-embracing process. Thus, if top-down determination applies to any processes (notably including the perceptual ones) then it applies universally, because there *is* only one process. Let us call this the *universe-process*.¹⁰¹

¹⁰¹ I will address in section 7.3 below the question whether the universe-process should be thought to embrace the *future* as well as the past and present.

I have assumed throughout that everything that exists is broadly physical. Adding this physicalist claim to the claims that (a) processes are metaphysically prior to their constituents (top-down determination) and (b) there is only a single process results in what I called *process monism*. As Schaffer defines it, monism is "the conjunction of the numerical thesis that there is exactly one basic object with the holistic thesis that the cosmos is basic" (Schaffer 2010: 42). Inasmuch as the priority ordering of what is more or less basic concerns *what determines what*, such monism implies that the nature of every individual thing in the universe (however we carve it up) is ultimately determined by the nature of the universe as a whole. In other words, every part of the universe – every object, let's say – has its nature in virtue of the position it occupies within the vast web of causal relations that is the entire universe.

In his defence of monism, Schaffer is anxious to overcome historical objections which insist that the monist is committed, absurdly, to their being just one thing – the whole universe. This motivates him to describe his thesis as what he calls 'priority' monism – namely the view that there are many different objects, but all are metaphysically subordinate or posterior to the one fundamental object, which is the universe as a whole (Schaffer 2010: 46). Common sense demands some admission of this sort – after all, our very talk of existence is founded on its application to the plurality of more or less familiar objects that occupy our environment. Certainly, we should allow that individual parts of the universe, however we carve it up, exist in some sense, even if we have cause to reject the notion that they exist absolutely (i.e. have a fixed nature at any given time).¹⁰² This has some implications for our notions of object identity which I will return to later in this section.

First, though, there is a more immediate worry whether experience-dependence can in principle be squared with process monism, or at least squared in a way that preserves the

¹⁰² We might be tempted here to avail ourselves of some convenient etymology. According to the Oxford English Dictionary, the origin of 'exist' is the Latin 'existere' or 'exsistere' meaning 'to stand out, be perceptible, hence to exist'. This is rather apt for our account of worldly objects, given that a metaphysics which treats processes as prior to their parts requires us to treat 'objects' as abstracted from the wider processes in which they participate. Objects, in other words, exist insofar as they stand out for us in perception. We must be wary, however, that this use of 'exist' does not sit entirely easily with our ordinary usage. In particular, it would be peculiar to say in ordinary speech that objects *cease to exist* when unperceived, but that is what the special use implies. A defender of the special sense of 'exist' would have to emphasise that the failure of objects to 'stand out' when unperceived does not mean that they, as it were, leave a hole in the universe. Rather, what it is to *be* an object in first place is to stand out, or be picked out, perceptually from its place in the whole universe-process.

natural view that worldly things themselves are experience-dependent. The worry is this: in treating the whole universe as fundamental in this way, it would be natural to assume that the universe as a whole has a fixed, absolute nature which fixes the nature of all of its constituents. This cannot be correct, however, as it conflicts with the perceptual phenomenon that led us to top-down determination in the first place, namely the experience-dependence of colour. This shows that objects' colour cannot be fixed by the antecedent nature of the whole universe any more than it can by the antecedent nature of some microphysical ultimates (as an atomist might have it).¹⁰³

The experience-dependence of colour – and in particular intersubjective colour variation – therefore demands a different conception of metaphysical relativism: one in which an object at any given time does not have an absolute nature fixed by how that object sits within the whole process or relational structure that is the universe, but instead has a different and unique nature relative to every process in which it participates. In the context of process *P*, object *O* has nature *N*; in the context of process *P** o has nature *N**, and so on. In the perceptual case we might say that the colour swatch is simultaneously blue-to-me (blue-in-my-experience) and green-to-you.

Again, I will suspend for now the obvious question of object identity raised by this proposal, namely the question of what justifies our saying that it is the same object in both cases if they have different natures. I will offer some remarks on this towards the end of the section. In the meantime we are sufficiently challenged to make sense of how this process-relativity could work in principle.

It will help to note that the talk of an object having different natures relative to different processes is already misleading – specifically, we are misled by talk of *different processes*. The claim that led us to monism is the claim that the universe as a whole, from

¹⁰³ This avoids a further epistemological worry that would arise were this sort of holism true. Imagine a world in which atomism was true and the nature of everything, including the whole universe, was determined by the antecedent nature of the microphysical atoms. We could then conceive of a world with all the same observable structure but in which the direction of metaphysical determination was reversed, i.e. the nature of every finite thing was determined by the nature of the universe taken as a whole. The two worlds would be empirically indistinguishable. Both would exhibit the sort of systematic structure that renders it amenable to bottom-up physical *explanation*, whether or not this structure were *metaphysically determined* from the bottom-up. The same set of empirically adequate bottom-up laws would apply to both worlds, but only in the atomist world would these laws carry any metaphysical significance.

its origin to the present, is a single process.¹⁰⁴ This seems empirically plausible and also avoids potential puzzles about individuating distinct processes. However, it there *is* only one process, then it cannot be correct to say that any object has different natures relative to *different* processes.

When we return specifically to perception and perceptual processes, the tendency to talk of different processes has its counterpart in our tendency to talk of *different experiences*, whether enjoyed by different subjects or by the same subject at different times. It is the different nature of these experiences that is supposed to account for the differences between colours as they are for different subjects.

So the contradiction is this: on the one hand, the subject-relativity of colour, etc., concerns a contrast between different experiences, experiences that we identify with *distinct processes*; but, on the other hand, process monism says there is *only one process*. On that view, what appear to be distinct processes are really arbitrarily demarcated regions of the one universe-process.

If process monism is correct, there must therefore be something wrong with claiming that distinct experiences are wholly distinct processes. The problem might begin with the very notion of 'distinct' experiences, which stands in need of clarification. With the simple metaphysical picture I have rejected the notion experiences as distinctively 'mental' entities, states or events that share the world with the more familiar objects, states and events that we know about and perceive. Viewing experiences instead as relations among worldly things (including, *inter alia*, subjects and objects of awareness) places limits on what we can mean when we talk of experiences being 'distinct'. For one thing, given the simple metaphysical picture, different subjects' experiences might 'overlap' in terms of sharing some constituents, as when the subjects are viewing the same objects at the same time.

What sense, then, can we make of talk of 'distinct experiences' within the context of a process monism? If there is fundamentally just one process, then my experience and yours, not to mention my experience at time t_1 and my experience at a later time t_2 , belong

¹⁰⁴ More precisely, what got us from the claimed universe-process to monism were the further claims of experience-dependence and top-down determination On its own, the notion of a single universe-process is perfectly consistent with atomism.

to that same process (the universe-process). 'Belong to' is of course too weak – if there is ultimately just one process then all these 'different experiences' are really one and the same. To reinforce this conclusion, we might note that the open-endedness of processes makes it difficult or impossible to demarcate any finite portions of the universe-process that could comprise our different experiences. It is therefore not obvious how we could non-arbitrarily place limits on the temporal and spatial extent of any perceptual process, assuming the causal interconnectedness of everything entailed by process monism.

Hence, ontologically speaking, all experiences are states of the whole universeprocess. Clearly, then, our ontological claim misses whatever it is that would allow us to make sense of our talk of distinct experiences. What might distinguish experiences if not their ontology? If my experience and yours do not differ in what they comprise (namely a whole state of the world) then they do self-evidently differ in respect of the *perspective* that they confer on the world. 'Perspective' here doesn't signal merely a spatial relation to the rest of the world (for even insentient things have that), but precisely the relation of awareness that constitutes experience. Note that the perspective provided by the subject obviously comes from within the world itself, so is in that sense an 'internal' perspective.¹⁰⁵ We might say that an experience is the world as it is from a certain perspective; or, better still, the world centred on the perspective provided by the physical subject. Different experiences involve different perspectives of awareness on the same universe-process. Put more prosaically, this amounts to the truism that different experiences involve different perspectives on the world.

Recasting the claim in terms of processes, we might now reappraise our starting claim that object O has nature N relative to process P, while it has nature N^* relative to process P^* . Granted that processes P and P^* overlap and are therefore continuous (this being the sort of observation that has pointed us towards process monism), the distinction between Pand P^* is seen to amount to a more or less arbitrary difference in which portion of the whole universe-process we choose to demarcate as the context for determining O. Object O therefore does not depend on *which* process it participates in, but on which perspective

¹⁰⁵ We have no obvious reason to suppose there is an *external* perspective – a 'God's-eye view' as it were. This perhaps bears on the question whether there is some overarching 'way the world is' that determines how it is from any particular perspective. A positive answer here might leave us wondering again whether this way the world is – the nature of the universe-process as a whole – ought to fix the nature of every one of its constituents absolutely, a result which would be at odds with experience-dependence. Experience-dependence would seem to suggest that the world is as many ways as there are perspectives within it.
we adopt within the whole universe-process when assessing its nature. This adoption of a perspective is of course precisely what we are doing when we perceive.

The perspectival role that subjects (and presumably other sentient creatures) play is just the familiar one of seeing things, hearing things, and so on. What remains puzzling, in light of the evident experience-dependence of colour, etc., is that shifting between different perspectives changes not only how things in the world appear but also how they *are*. The world as it is for me (to my perspective) may be different to the world as it is for you. Given my monistic appeal to a single universe-process, we cannot allow that these are two different worlds. Indeed, we can easily show that if there were 'two worlds' then they would at least overlap. After all, you can feature in my experience as a physical subject, and *vice versa*. What we have is not two worlds but one world, as perceived from different perspectives within that world.

If worldly things are truly different to different perspectives (blue to me, green to you, and so on), where does this leave the status of our factual claims about the world? Claims about objects' colour, for example, are relativized to the subject, not because there is some inaccessible fact of the matter for which we must substitute a mere appearance, but because objects' colours just are subject-relative.¹⁰⁶ As explained in section 7.2.4, we are in a slightly different position in respect of so-called 'primary' qualities, but this is arguably not because there is an absolute fact of the matter about, e.g., shape and mass, in a way that there isn't about colour. For one thing, we can understand our claims about such properties to be already implicitly or explicitly relativized, not directly to the subject, but to measuring tools and standards which thereby act as a sort of intermediary between the objects and us. This is enough to grant them a degree of independence from variation among individual subjects, enough to deliver an ordinary sort of objectivity.

If the same object has different properties relative to different subjects, and yet these properties (as in the case of colour) are genuine qualities inhering in the object itself, what entitles us to say that we see the *same* object in such cases? Here, the intersubjectivity that presents the puzzle also provides a solution: quite simply, we can *see* that we are both looking at the same object. I can see the object and I can see you looking at the same

¹⁰⁶ Plainly, not all of our claims about objects' properties need to be relativized to the subject in this way, notably those concerning 'primary' qualities. The reasons for this were discussed in section 7.2.4 above.

object. Furthermore, and notwithstanding some potential disagreement about colour, enough of our respective descriptions of the object tally to affirm that we are indeed seeing the same thing.

This rather pragmatic answer, however, might not overcome a related worry: if objects lack any absolute properties, how can we be confident that there are any real commonalities between the experiences of different subjects, even with respect to, say, size and shape? The worry might be that there is nothing to 'anchor' such commonalities and so justify our assumption that different subjects' experiences match in at least some key respects. For example, we can imagine that the colours might not just vary a little but be wholly permuted (e.g. inverted) for different subjects, in a way that broadly preserves the structure of similarity relations between them and so supports use of a common colour language.

Again, intersubjectivity is the key to a solution. *Ex hypothesi*, the nature (e.g. colour) of an object is determined by the nature of the process in which it is perceived. Now, the nature of the process by which *you* see an object *O* is (potentially) accessible to *me*. I am, or can become, aware of the elements of that process – the surface structure of the object, the spectral profile of light reflected from it, the way your visual system responds, and so on. That alone is no comfort: the same process might, after all, appear very different to you, in which case we are back where we started. Crucially, however, it is open to me to inspect *my own* visual process 'from the outside'. From inspecting your perceptual process, I already know the relevant properties of the object and reflected light, and I can now investigate my own visual system's response in the same way I investigated yours (the mechanics of which investigation I delegate to my neuropsychologist friend). In the event that your physiological responses to the same worldly features (the object's shape, colour, etc.) prove closely similar to mine, I will have some justification for taking the object to appear (indeed to *be*) much the same for you as it is for me.

The puzzle of how we establish the identity of objects that might have different properties to different subjects potentially gives way to a deeper puzzle: what can we say about objects that are unperceived? Are apples green when no-one is looking at them? Does an apple even *exist* unperceived? It would be an offence to intuition to say 'no' to either question. However, at least in respect of colour, the answer can't quite be 'yes' either if my metaphysical account is correct. According to that account, an object's colour is determined by the awareness relation or process of which that object is part. To the perspective of a different subject the apple may therefore have a different colour. In the *absence* of a perspective, there is no process in place that will determine what the apple's colour is. The colour is, in that sense, undetermined. But this is not to say that the apple is positively *colourless*.

We might think that the object's existence is similarly undetermined, especially if we take the object to be nothing more or less than its qualities. That being so, it would be similarly untrue to say that an unperceived object positively does *not* exist – that the relevant part of the universe-process positively *lacks* all the qualities we find when we have a perspective of awareness upon it. Here, as with the question of object identity, we might look for a pragmatic basis for our ordinary talk of things existing when no-one is looking at them – one which acknowledges our ability to resume something like our previous perspectives on the world and so bring the same objects back into awareness.

7.3 The presence of the past: the temporal structure of experience

At the beginning of this chapter I explained that using the simple metaphysical picture in order to accommodate various puzzling perceptual phenomena to the natural view forces us to embrace two main metaphysical consequences. The first of these was the subjectrelativity of certain perceived qualities, which I have discussed in section 7.2 above. The second concerns the temporal structure of experience, and in particular the 'presence of the past' in perception.

The proposed explanation of such phenomena as hallucination and cognitive penetration of perception drew on the simple metaphysical picture to claim that such experiences involve direct awareness of past objects, made possible by our memory of those objects. The current experience, be it perceptual or hallucinatory, restores an appropriate causal link between the subject and some past object(s), where 'appropriateness' might include amongst other things the fact that the causal link is (historically if not currently) mediated by the sense organs. Such phenomena, and the proposed explanation, force us to make a striking claim about the temporal structure of experience, namely that we can be aware of objects and events that belong to the past. Of course, such a claim is already forced on us by the natural view of perfectly ordinary perception: given that my seeing an object in my environment is always mediated by a causal process, it follows that *every* perceptual experience is awareness of something as it *was* rather than as it *is*. As noted earlier, the time gap involved is most often negligible, but obviously increases with distance, so that I see the moon as it was eight seconds ago, and the Andromeda galaxy as it was 2.5 million years ago.

If it is tempting to talk here of 'the presence of the past', that is because we can exploit an ambiguity in 'present' and 'presence', as we use these words to describe our awareness and the things we are aware of. On the one hand, we talk about things in our environment being present to us in awareness, despite the fact that what is thus presented is inevitably past. On the other hand, those events in our environment that occur in the present – which, for the sake of argument, we might specify in terms of simultaneity with some current state of our brain – are themselves not (or not yet) present in our awareness.

We might therefore distinguish two senses of 'presence': an *experiential* presence and a *chronological* presence. An object is present in the experiential sense if the subject is directly aware of it (where by 'direct' here I mean that the subject is not aware of the object only via awareness of something else, like an image of the object). An object – an event, say – is present in the chronological sense if it is simultaneous with the current state of the subject. More generally, we might think of chronological *co-presence* as a relation of simultaneity holding between two or more objects, where one may or may not be a physical subject. We have ways of measuring chronological co-presence or simultaneity in terms of causal separation and using certain standards to calibrate this separation (being a constant and, so far as we know, a physical limit, the most useful standard here is the speed of light).¹⁰⁷ Crucially, the causal separation between perceived worldly objects and the physical perceiver is such that nothing is, as perceived, ever quite present in the chronologically present.¹⁰⁸

¹⁰⁷ Given Einsteinian relativity, we can specify temporal separation by reference to causation only within a given inertial reference frame, so that chronological measures of simultaneity and non-simultaneity admit of no absolute values.

¹⁰⁸ Of course, many objects that are experientially present will continue to exist through whatever period is required for the object to perceptually affect the subject, in which case they will still be chronologically present, just not *as perceived*, i.e. the chronological location of the object at the moment it is perceived is always earlier than the chronological location of the subject at the moment she perceives (assuming some spatial separation between object and subject).

This inevitable time-lag in perception has often been seen as ruling out direct awareness, although this only follows from an antecedent commitment to the effect view of perceptual experience, i.e. the notion that perceptual awareness involves the having of an 'experience' understood as a state or event that exists only in the present (see Power 2010: 97-8; for a statement of the 'time-lag argument' against naïve realism, see Robinson 1994: 80-1). The natural view and simple metaphysical picture simply reject this notion of experience: instead, they claim that the 'experience', as such, is nothing more or less than the temporally extended perceptual relation or process itself.

Awareness of the past is thus effectively built into the natural view and simple metaphysical picture from the outset. My concern in the following two sections is therefore not to defend the claim that we can be directly aware of the past, but to tease out the likely consequences of accepting this claim. Here again, we can see our task as a positive one: rather than seeing the natural view as generating a *problem* around temporality, we can choose to see it as licensing our use of perception as a straightforward source of *data* about the nature of time. To that end, I will outline what we might learn about the temporal structure of the world by taking experience at face value.

In the next section I will consider the extent to which the natural view and its insistence on the 'presence of the past' can be squared with the main metaphysical views of time, namely presentism and four-dimensionalism. Then, in section 7.3.2, I will examine the consequences of combining the presence of the past with the experience-dependence of worldly things and their qualities.

7.3.1 What does experience tell us about time?

It will help us tease out the implications of the natural view for our picture of time if we assess its fit with three main metaphysical theories of time, named *presentism*, *eternalism*, and the *growing universe theory* (see Markosian 2014: s.6). These differ according to the status they grant to the past, present and future. According to presentism, only the (chronological) present exists. Eternalists, meanwhile, claim that everything – past, present and future – exists equally. Eternalism is a form of what is known as *four-dimensionalism*, since it treats objects as fundamentally extended in time as well as

through the three dimensions of space (Rea 2005). It also appears in the guise of the *block universe theory*, according to which the whole history of the universe (what appears to us as past, present and future) is an eternally-existing 'block' with no privileged 'now' (Dowden 2013: s.8a). Besides eternalism, there is a further version of four-dimensionalism which provides our third main theory of time. This is the *growing block* or *growing universe* theory, which holds that only the past and present exist, but not the future. Since the present is forever changing, to become the past, the theory entails that the universe is constantly 'growing'.

Now, of presentism, eternalism and the growing universe theory, we can quickly see that presentism is at odds with the natural view. The natural view allows that we can see the past, so presentism is ruled out straight away since it requires that only what is chronologically present exists (Power 2010: 112).¹⁰⁹ This move effectively reverses the traditional time-lag argument against direct perception. The time-lag argument assumes the effect view of perceptual experience, namely the view that perceptual awareness is a state or event caused by, and so temporally later than, the object perceived. Granted the effect view, it follows from the time-lag inevitably involved in perceiving our environment that the experience occurs after the causative state of its object has ceased to exist. The conclusion of the time-lag argument is therefore that the object is no part of the experience and direct perception of the sort proposed in the natural view must be false (Robinson 1994: 80-1). However, with the rejection of the effect view in favour of the simple metaphysical picture (see chapter 2 above), the time-lag argument fails and we are furthermore moved to reject presentism since we can perfectly well *see* the past (see Power 2010: 97-8).

We might therefore think that we are thereby committed explicitly to a fourdimensionalism of either the eternalist or growing universe varieties. After all, I have already suggested that accounting for the experience-dependence of certain perceptible qualities, in light of the natural view and simple metaphysical picture, leads us to a conception of processes as fundamental, rather than the more or less momentary parts that

¹⁰⁹ Presentism is in any case challenged by special relativity, according to which there is no absolute simultaneity but only simultaneity relative to an arbitrarily-chosen inertial reference frame (see Sider 2001: 42 ff.). This has implications also for the growing universe theory since it privileges a certain set of simultaneous events (the present ones) as the 'leading edge' of the growing universe (Miller 2013: 352 ff.). For a defence of the 'moving spotlight' form of eternalism from objections appealing to relativity, see Skow (2009).

they comprise. Indeed, I further suggested that this top-down metaphysical picture lends itself to a monistic ontology in which the whole universe-process is fundamental. Earlier I hedged my bets around the question whether this whole universe-process should be considered to comprise the whole history of the universe up to the present, or whether this should be extended to include the future as well. The difference would of course be precisely that between eternalism and a growing universe theory.

As it is, neither view appears to fit our metaphysical account terribly well. Eternalism in particular fails to acknowledge the central phenomenological datum that we might describe as the evident passage of time (Miller 2013: 353). Its insistence that past, present and future are equally real is glaringly at odds with our intuitive sense of change – the sense of a *now* that gives meaning to the notion that things (events, say) which are chronologically present become chronologically past (rather than things' merely occupying earlier and later stages of an eternal, four-dimensional universe).¹¹⁰

This might seem a good reason to favour the growing universe theory, for that at least seeks to capture our sense of a changing present through the notion of history's ever-expanding 'edge'. Here too there are problems, however. The growing universe theory, like any form of four-dimensionalism, takes every moment of the world's history to be equally real, but this sits almost as uncomfortably with intuition as eternalism's insistence that such reality extends also to the future. Some moments in the past include my five-year-old self's coming to be aware of things. But this no longer seems as real as my coming to be aware of things *now*. While eschewing the reality of things yet to happen, the growing universe theory seems unable to express this intuitive difference between my current experience and that of my five-year-old self.

How, then, to square the following claims I have proposed on the back of the natural view and the simple metaphysical picture?

a) Experience (one's perceiving X) is a perspectivally-centred state of the whole universe-process (the whole world) from its beginning to the present.

¹¹⁰ The difference is that famously captured by McTaggart's distinction between dynamic and static conceptions of time that he refers to as the 'A series' and 'B series' respectively (McTaggart 1908). Eternalists may try to recapture this sense of time's dynamism by appeal to some special role of the present as a 'moving spotlight' roving over successive parts of the four-dimensional block universe (see Skow 2009).

- b) The universe-process is all there is and it is fundamental.
- c) A subject's current experience is in some sense more real than her past experience.

There is no logical inconsistency here, but the apparent contradiction arises because (a) and (b) seem to assert the equal reality of everything past and present, while (c) seems to demand some special reality for current over past experience. The seeming contradiction is overcome once we recognise that past and present experiences of the same subject are not two separate 'things' but different relations of awareness holding between the subject and the rest of the world. Although my five-year-old self's experience was, at the time I was enjoying it, a perspectival state of the whole world, it is no *longer* that, as I recall it *now* from a vantage point decades later. *Now* it is my very experience of *recollection* that constitutes the world as it is from my perspective. As such, my five-yearold self's experience is really no longer an experience at all. Understood as a process or physical relation, it has now become part of the larger process that is my act of recollection.

It is important to note that my past experience has not, as it were, 'disappeared' or ceased to exist, leaving only a (memory) trace. That is a view we might attribute to a presentist, but it is not one that fits with the simple metaphysical picture of what is going on when we recall a past experience (or, more accurately, recall an object previously perceived). According to the simple metaphysical picture, the perceptual relation in which my five-year-old self stood to that past object is still 'there', insofar as it is an essential part of my current recollection of that act of perceiving. In other words, that historical perceptual process is now part of the larger process of recollection in which I am currently engaged. This process of recollected objects. Given top-down determination, it follows that the earlier perceptual process is now subordinated to the larger process that is my experience of recollection. It is in this sense that (c) is true – that my current experience is more 'real' than my past experience. That is, my current experience is constituted by a larger process (indeed, the whole universe-process if process monism is true) of which the

earlier perception is but a part. My current experience is therefore the more fundamental process.

We therefore have a way to make sense of the difference between current and past experience: current experience is, granted process monism, the complete and therefore fundamental universe-process. Current experience is, in this sense, *what there is*. When I think of one of my past experiences, that past experience – as a process culminating in some historical state of my brain – is no longer the complete, fundamental universe-process. The process has continued subsequently, and indeed is now what we might consider a *new*, current experience – namely the experience of recollecting that earlier, perceptual experience.

Note that, although there is a sense in which a current experience becomes past and so less 'real', the same is not necessarily true of the *objects* of our experiences. As we go from seeing an object to visually recalling it – as the object recedes into the chronological past – it nonetheless remains *experientially* present. Of course, the subject's *perspective* on the objects has changed between seeing them way back when and recalling them now. The shift in perspective is precisely that from *seeing X* to *recalling X*.

It remains unclear whether the experiential presence of the past lends support to, or is even compatible with, the growing universe theory. The growing universe theory grants a special role to the ever-advancing present, each moment of which continues to exist even as it is superseded. Meanwhile, the picture that seems to fall out of the natural view is one in which the chronological present plays only the slenderest role. Everything that features in our perceptual experience is already chronologically past, if only by the tiniest margin in most cases, so the chronologically present is not something we ever actually encounter. Reality, so far as we perceive it, is all in the past.

Although perception makes us aware of changes in worldly objects and so enables us to form the notion of chronological change, it also weakens any deep metaphysical distinction between past and present, since both are equally present in the experiential sense. What changes over time, for us as subjects, is not that our awareness inhabits a different time-slice of the universe, but that our awareness amounts to a different perspective on – or within – the universe.

7.3.2 The experience-dependence of the past

A more obvious conflict with the growing universe theory, and four-dimensionalism in general, comes once we ally the presence of the past with a top-down metaphysics. Fourdimensionalism assumes that the past, once past, has a fixed nature for all time. If, however, top-down determination is the fundamental structuring principle then the past becomes metaphysically dependent on its relation to the present, as processes leading up to time t_1 give way to, and so become absorbed within, wider processes leading up to time t_2 . This leaves open the possibility that the past is not, after all, fixed for ever, because, once past, its nature becomes metaphysically dependent on its place with the larger processes in which it is subsequently embedded (indeed, the one evolving universe-process, if process monism is true).

To see how this works, and to emphasise the temporal dimension, let us consider the perception of an *event*. The upshot of applying our top-down metaphysics is that the nature of that perceived event for the perceiving subject is determined only once that event has participated in the perceptual process whose 'conclusion' comes later, chronologically speaking (the 'conclusion' here we might stipulate to be the neural state of the subject that issues most immediately in acts of recognition etc., even though the process carries on indefinitely). So, although the subject might see an event that happens at time t_1 , its nature (to the given subject) *at time* t_1 is only determined upon the 'completion' of the perceptual process at time t_2 . This yields the seemingly paradoxical result that the event is, as veridically perceived, *F* at t_1 but only comes to possess that quality *F* in virtue of its participation in a process that is not complete until t_2 . The seeming paradox is this: we see the object as it was at t_1 , namely *F*, but it didn't come to *be F* until the experiential relation came to fruition at t_2 . So it couldn't have been *F* at t_1 at all.

We might be tempted to say that what is involved here is the determination of the past by the present. But, once again, we must be wary of the latent ambiguity in our use of 'present'. What is fundamental, according to the top-down metaphysics, is the whole experiential relation or process. The determining role is therefore played not by what is chronologically present (e.g. the current state of the subject's brain) but by the whole experience, including everything that is *experientially* present. And that, most obviously, includes the (chronologically past) object of awareness. The perceived event is experientially present to me, even although it is chronologically not simultaneous with me at the moment I become aware of it. The crucial point is this: although we might correctly say that I become aware of the event at time t_2 , even though the event itself occurred at time t_1 , it does not follow that the *experience* happens at time t_2 . The experience, according to the simple metaphysical picture, is itself a temporally extended process incorporating, *inter alia*, both the event at t_1 and my perceptual brain state at t_2 .¹¹¹ Granted this, it is not the case that the event has its nature determined specifically by some later event (i.e., whichever event might be said to 'complete' the perceptual process). Rather it has its nature determined by the whole process of which those (chronologically) earlier and later events are discernible parts.

The seeming paradox can thus be resolved. It is not the case that the object only became F at t_2 . It was already F at t_1 , just as it appeared perceptually. What it *is* for the object to be F at t_1 is precisely for it to stand in a certain relation to the subject in her current state at t_2 .

Of course, even if we can strip it of the charge of contradiction or paradox, we should acknowledge that this picture of past events being determined by their involvement in processes that extend to the present remains highly counter-intuitive, if not downright perplexing. Any residual puzzlement stems not just from the claim that objects have their nature determined 'retrospectively', as it were, but that this nature could be different, not only for different subjects, but also for any given subject at different times. This suggests that an object's nature at a certain time t – even its nature relativised to a given subject – is not fixed but might change as it goes from being perceived to being recollected.

One thing to say is that this is not obviously a problem in practice – we seem capable enough of recalling objects accurately in respect of their perceptible qualities. Of course, we frequently recall things poorly or inaccurately, but the inadequacy here might typically be attributable to one or both of two factors: firstly, the disruption of the causal link

¹¹¹ As noted previously, this means that 'experience', in the sense I have used it, i.e. to denote the relation of awareness involving subject and object, is not to be understood as synonymous with 'act of awareness'. Consider glancing at the sun: here the relation of awareness is describable as a process whose most relevant part – in which light emitted from the sun reaches the subject's retina – lasts some eight minutes. To say that the experience lasts eight minutes would be quite misleading if we were using 'experience' synonymously with 'act of awareness', for the latter in this case amounts to a mere glance.

between the subject and remembered object that underwrites the memory awareness; and, second, confusion of the target object of memory with one or more other objects previously perceived, such that these latter objects are in fact the actual objects of awareness. It is therefore not obvious that, taking our experiences at face value, we have grounds for thinking that objects change in respect of their perceptible qualities as they go from being perceived to being recalled.

Nonetheless, top-down determination would seem to leave open the possibility that, for example, the green flash I saw in 1979 is now blue for me as I recall it in 2014, perhaps because my visual system has been altered in the interim, and not because my memory awareness is confused or otherwise inaccurate. In truth, this possibility presents no principled challenge beyond that already posed by inter- and intrasubjective variation in colour perception. In all these cases, the object of awareness can possess different qualities according to the perspective from which the subject is aware of it. In cases we considered in chapter 6 the perspectival differences had to do variously with differences in the physiology of different subjects, or changes in the background conditions of perception for any individual subject. In the current hypothetical case, in which the 1979 flash of light was green for me in 1979 but blue for me in 2014, the perspectival change is a temporal one, and the change in the flash's colour might, for the sake of argument, reflect (but not be *caused* by) a change in how my visual system processes visual memory traces.¹¹²

Where does all this leave our debate, discussed in the previous section, concerning the various competing metaphysical views on time? I noted there that the metaphysical subordination of past objects and events to the wider processes they give way to (and perhaps ultimately to the universe-process as a whole) means that the top-down metaphysics is not likely to prove consistent with the growing block theory any more than it is with eternalism, since both require that all objects or events at all stages of history are equally real and retain their nature through the passage of time. The alternative picture presented here is one in which the past may be very much present in the experiential sense, but in which past objects have their nature in virtue of the (perceptual or memory) relation they stand in to current subjects. As no object has, on this view, an absolute nature independently of how it exists to one or more perceivers, there is no such absolute nature

¹¹² Any dispute about the 'actual' colour of the object (the flash, in this case) might be settled on pragmatic grounds, in much the same way as we might settle competing colour claims *between* subjects (see section 7.2.5 above).

to remain fixed through the world's evolution. Rather, as the world evolves, the status of that object (that moment in history) within the world evolves correspondingly. It exists to the same or different subjects in diverse and different ways, being perceived from diverse perspectives and remembered from others.

Top-down determination therefore yields neither presentism, eternalism nor the growing universe theory. How, then, does it leave our understanding of the temporal nature of the world? If it is tied to a process monism as I have suggested, it holds that what there is at any (experiential) moment is the universe-process as a whole, from its beginning to the present. Within that whole universe-process, certain features (objects, events, qualities) stand out according to the perspective that is provided within the universe-process by a physical subject. There is some sense, then, in which all of the past is potentially 'there' awaiting an opportunity to stand out to some perspective. Having said that, there is little or nothing we can say about the past unless it does stand out, and if the view offered here is correct this is not merely an epistemic problem but a metaphysical one – it is only by standing out in the perspective of a subject that objects come to have any determinate (if not absolute) qualities; qualities, indeed, by which we might mark out those objects' very existence.

Not only is the nature of the past not fixed, but much of what is past is presumably lost to us forever and will never stand out again for any subject. This will be so if causal linkages between then and now fail to preserve information about certain states of the past. In that case, we will no longer be able either to *perceive* those elements of the past in the way that, for example, we perceive the Andromeda galaxy as it was 2.5 million years ago. Nor will we be able to remember them. Such 'lost' elements of the past will no longer stand out for our awareness, although they leave no 'hole' in the universe.

7.4 Conclusion - metaphysical implications of the natural view

My purpose in this chapter was to outline the likely metaphysical and ontological implications of the natural view. In assuming the truth of the natural view I was able to treat experience as a straightforward source of data or evidence about the nature of the world. Some of the claimed implications of the natural view are bound to be highly controversial – not only experience-dependence of worldly objects and qualities, but also

the process monism that I suggested might follow. I sought, however, to cash out experience-dependence via a broadly physicalist ontology, resulting in what I called a topdown metaphysics in which wholes determine their parts rather than *vice versa*. This allows us to account for experience-dependence in a way that avoids idealism, as befits a theory that eschews talk of 'minds' and 'mental states' from the outset.

Key features of the resulting view are as follows:

- The universe is governed by a basic metaphysical structuring principle of topdown determination, by which the nature of whole processes determines the nature of their parts.
- Assuming that the universe is a single process, top-down determination entails a process monism in which the whole world from its beginning to the present is fundamental.
- A top-down metaphysics points to a conception of time that admits the experiential presence of what is chronologically past, but which (*contra* four-dimensionalism) does not support the notion that past and present experiences are all equally 'real'.
- None of presentism, eternalism or the growing universe theory fits the picture of time bequeathed by a top-down metaphysics. In particular, the requirement of eternalism and the growing universe theory that a given object's nature at time *t* is settled once and for all time is inconsistent with the observed experience-dependence of worldly objects and qualities, and with the claim that the past is determined by its role within processes extending to the present.

Plainly, the resulting picture is, in at least some respects, highly counterintuitive. In particular, the notion that the past is not fixed, but remains dependent on its participation within the ever-evolving universe-process, is decidedly hard to fathom, even if we can rid it of contradiction. All the same, the core insight that the past is (experientially) present is arguably quite intuitive. Our recollections do seem to 'reach into the past', while an understanding of the time gaps involved in perception doesn't necessarily override our sense that we really do *see* distant celestial objects. The most puzzling aspect of

temporality, as revealed in perception, therefore results from the claim of top-down determination rather than the presence of the past.

CHAPTER 8 - CONCLUSIONS

The aim of this thesis was to defend what I have called a natural view of perceptual experience, according to which perception is just what it seems to be: a straightforward encounter with worldly objects and their qualities. In seeking to flesh out what such a 'straightforward encounter' is meant to entail, I explicitly disavowed all talk of distinctively 'mental' states or events, believing there to be nothing of the sort, or at least nothing substantial deserving the name. Instead I offered a metaphysically simple, if not austere, account of perception in which the 'experience', if we must use that potentially misleading term, is nothing more or less than a certain sort of relationship among familiar, physical things – human organisms and the various things that make up their environment.

According to this *simple metaphysical picture*, perception involves nothing besides these familiar, physical things, as they are bound together in certain states of affairs that we can describe in physical terms as perceptual processes. These are the sort of processes long familiar from empirical research into the workings of the senses. Crucially, perception involves no distinctively mental state – 'experience' in a more substantial sense – as the final upshot of the perceptual process. This sets the simple metaphysical picture against what I called the *effect view* of perceptual experience. It also inoculates the view against the weirdness that attends the sort of claims often made by naïve realists, to the effect that worldly objects are constituents or ingredients of our perceptual experiences.

Having established that neither conceptual nor empirical considerations decide in favour of the effect view over the natural view, I went on to address some of the more puzzling features of sensory awareness, features that have widely been thought to undermine the natural view. Chief among these are perceptual illusion and hallucination. I described how proponents of contemporary naïve realism have sought to defend something like the natural view from a certain kind of argument that trades on the possibility of such seemingly aberrant sensory phenomena. This kind of argument seeks to generalise from some alleged failure of sensory engagement with the world in the 'bad' case (illusion or hallucination) to the conclusion that veridical perceptual experience must also fail to engage us with the world in quite the direct way claimed by the natural view. Naïve realists have shown that we can readily dismiss this argument in respect of many illusions, essentially by pointing out that we can be misled by how things appear even though we are directly aware of them in the way the natural view requires.

Hallucinations prove, as a class, less easily accommodated to the natural view. The standard naïve realist defence to the argument from hallucination has invoked the disjunctive analysis of appearance statements, whereby a given statement about how things appear might be made true by quite different states of affairs – one a genuinely perceptual encounter with worldly things, and the other a relation of awareness to something unworldly or uninstantiated, or else not really a relation of awareness at all. I noted that such disjunctivism renounces the virtues of a unified explanation of what appears to be potentially shared phenomenology across perceptual and hallucinatory experiences. This is the unified explanation offered by conjunctivist or common-factor theories, which ascribe the potentially similar phenomenology of perceptual and hallucinatory experiences to the occurrence of a single kind of state or event – the experience itself – in both cases. However, it was with hallucination that I was first able to trumpet the explanatory power of the simple metaphysical picture, as a means of giving a unified explanation of perceptual and hallucinatory phenomenology that nonetheless kept all that phenomenology 'out there'. In other words, the phenomenology even in hallucination is attributable to worldly objects – not necessarily objects currently before the subject, but objects previously perceived and still potentially available to the subject's awareness by virtue of a causalinformational link to some persistent effects on her brain ('memory traces').

I went on to show that this memory-based explanation of hallucination can act as a launch pad for similar claims about other features of perception that might otherwise appear to make trouble for the natural view. Most notably, this includes those states of awareness in which past experience, beliefs or expectations evidently play a role in determining the way things appear in a robust, phenomenal sense. Here, the role of such cognitive factors was explained as a kind of partial, simultaneous awareness of one or more previously perceived objects while genuinely perceiving some current object. So, when I see a familiar face, what contributes to how it looks is not just awareness of that face before me now, but also some or perhaps every previous moment of awareness of that same face, in a sort of 'layering' of the past upon the present in a way that is seamless and so not obvious. I showed that this allows the simple metaphysical picture to accommodate the generality of perception as well as its immediacy.

As far as this step in the thesis, I was able to put the simple metaphysical picture in the service of naïve realism, that is the conjunction of the natural view with the *realism claim*, namely the claim that the objects and qualities of which we are aware exist and have their nature independently of their being perceived. The decisive twist in the overall project then came with consideration of certain features of perception – veridical perception, no less – which prove incompatible with either the natural view, the realism claim, or both. These features of perception are ones which reveal some subject relativity or experience dependence of perceptible qualities, and I described aspects of colour perception and indeed colour itself which exhibit this experience dependence in various ways.

This experience-dependence of at least some of the phenomenal qualities which seem to inhere in worldly objects plainly rules out naïve realism if that theory is understood as the conjunction of the natural view and the realism claim. At this point, we might be tempted to cut our losses and retreat from the natural view, rueing the futile effort expended in trying to square it with hallucination and so on, but glad enough to walk away with realism intact. However, I noted that one could coherently preserve the natural view if one were willing to dispense with the realism claim. Crucially, I argued that simple metaphysical picture offers a way to make this gambit more than merely logically possible, but perhaps even metaphysically intelligible. The suggested solution involved viewing 'experience-dependence' through the prism of the simple metaphysical picture, the upshot being the claim that objects are dependent in respect of at least some of their properties (e.g. their colour) on the nature of the wider (perceptual) processes in which they participate.

At this point I suggested that this dependence relation between objects and the processes in which they participate – a relation I called *top-down determination* – might in fact not be a feature peculiar to some or all perceptual relations, but instead a wholly general feature of the world. The universal scope of top-down determination was defended chiefly on the grounds that it offered the virtue of a unified explanatory and metaphysical principle. At the same time, it was intended as speculative in the most positive and constructive sense, namely as the floating of an idea in order to test its explanatory power.

My task in chapter 7 was therefore twofold: firstly, to draw out the likely metaphysical implications of treating top-down determination as a universal metaphysical structuring

principle; and, second, to assess whether it might be consistent with, or able to explain, the way the world appears to us perceptually. As regards this second concern, it might seem that much of the work had already been done in chapters 3 to 5, and that I had in fact managed in that chapter to show that the natural view and simple metaphysical picture could explain various sensory experiences quite satisfactorily without recourse to top-down determination. This in fact presented a further challenge to the proposed metaphysical thesis based on top-down determination: not only to explain those phenomena which *don't* submit to naïve realism and its explanatory virtue (a virtue which we saw implies *bottom-up* rather than top-down determination, i.e. atomism rather than holism), but to explain why some properties of worldly objects *do appear* consistent with naïve realism (and the realism claim in particular) in spite of its claimed falsity. To that end, I sought to explain the apparent experience-*independence* of certain qualities – the so-called 'primary' qualities of size, shape and so on – as a consequence of the distinctive way in which subjects are related perceptually to those qualities.

The claims in chapter 7 were therefore intended to uphold the natural view of perception as acquainting us straightforwardly with the world around us, by showing that it could be defended not only in light of some seemingly puzzling sensory phenomena but also within the ambit of a wider metaphysical picture of the world that shares the natural view's virtues of simplicity and parsimony. While that might seem to present the challenge as one of fitting our wider metaphysics about a certain view of perception, I argued that the natural view itself, in eschewing talk of distinctively 'mental' states or events, implies something different and perhaps more positive about the metaphysical task. Instead of treating perception as a peculiar, local phenomenon – something tied to the sentient beings scattered around one or more corners of the universe – that makes disproportionately large metaphysical demands, we should see perception as a peculiar, local phenomenon that reveals the world as it is. In other words, it is not that we fashion our metaphysics in order to explain perception, and then apply that metaphysics to the world in general. Rather, if we take the world at face value, as the natural view recommends, it is the nature of the world itself, as *revealed* in perception, that guides our metaphysics.

Having said all that, one lesson we can learn from experience is that the world does not always reveal itself straightforwardly. Indeed, it is perhaps never revealed quite as straightforwardly as we might think, and indeed as the natural view might suggest. Sometimes we need to look very closely to work out what is going on. As cognitive penetration and perceptual learning reveal, even the most seemingly straightforward experiences prove to involve more than simple awareness of what is perceptibly before us.

One implication of this is that we are rarely, if ever, wholly present, in the sense that our awareness at any moment encompasses not only what is before us now, but also what was once before us but is now absent. To put it metaphorically, we perceive the present through the lens of the past, perhaps inescapably so. As I explained in chapter 7.3, combining this chronological 'depth' of awareness with the notion that we should take experience at face value has surprising implications for our understanding of time, implications brought out in the distinction between chronological presence and experiential presence.

What, finally, should we say about the mind? The view proposed here started from a narrow concern to explain the nature of perceptual experience, but proved capable of accounting for perhaps the full diversity of what we might call phenomenal consciousness more generally. Inasmuch as phenomenal consciousness is paradigmatic of our mental life, we might extend our conclusions about the nature of awareness to the mind more generally. I have said that our having an experience does not involve some internal and peculiarly 'mental' state or event, but is simply our being related to the world through awareness in a more or less straightforward way. By extension, one's having a mind, so to speak, does not signal one's possession of some repository or medium of 'experiences', but rather that one can and does stand in certain relations to the world. Those 'mental' relations would typically be described as ones of awareness and thought. My chief concern has been awareness, although I hinted in my account of cognitive penetration and perceptual learning that we might go a long way to explaining the nature of thought in terms of the simultaneous, 'folded' awareness of worldly things both past and present. Making good that suggestion is beyond the scope of this thesis, but it should be obvious at least how we might hope to cash out the notion of thinking in terms of one's standing in certain sorts of relations to worldly things, and without recourse to talk of 'thoughts' construed as inner states or events.

Since the simple metaphysical picture has no place for peculiarly mental states or events, all the onus for explaining the nature of phenomenal consciousness is placed on the world. Happily, those seemingly anomalous features of our subjective lives that did so much to motivate talk of 'minds' and 'experiences' in the first place – features like illusion, hallucination, dreams, and so on – prove to be explicable without our taking such talk literally as picking out metaphysically peculiar entities. However, passing the explanatory responsibility back to the world forces us to revise some well-entrenched assumptions about the nature of the world and our place within it. In particular, the failure to uphold the explanatory virtue of naïve realism – its promise to explain the nature of experience by reference to the antecedent nature of the worldly thins perceived – may be attributable ultimately to a mistaken adherence to an atomistic conception of the world in which big things and processes are 'built out of', and so determined by, littler things. I have suggested that adopting a holistic, and indeed monistic, conception instead might allow us to accommodate the seeming experience-dependence of certain qualities within a broadly physicalist framework.

To many, such a radical and wholesale change to our picture of the world will seem an unacceptably high price to pay for a unified account of phenomenal consciousness – an unreasonable shifting of the explanatory burden from the local matter of mind to the world in general. However, the natural view precisely offers no reason to think that perception reveals something merely local about ourselves – rather, it ostensibly reveals the nature of the world, as befits its purpose. Taking the world more at face value therefore allows us to treat it, not as something that lies *behind* our experiences, but as laid out before us *in* our experiences. This, if you like, is our 'phenomenological' reason for basing our metaphysical view of the world on the nature of awareness.

Still, it leaves the natural view to stand or fall according to the viability of some metaphysical scheme that can support it. I have suggested one such scheme here, and justified it on the grounds that it seems to fall quite 'naturally', as it were, out of the natural view, and that it appears both simple and parsimonious. However, further work will be needed to determine whether it is viable on independent grounds.

If its viability can be established, the metaphysical view proposed offers its own advantages beyond merely explaining the various features of phenomenal consciousness. Specifically, it straightforwardly avoids any metaphysical or ontological puzzles about the relationship between mind and world, puzzles which drive so much of philosophy past and present. Relatedly, it promises to avoid at least some of the epistemological worries that arise from viewing mind and world as in some way distinct, or from taking perception to be in some way less than 'direct'. I have said little explicitly about how the natural view and simple metaphysical picture impinge on some of these philosophical debates, but it is here that it is likely to have some of its greatest advantages.

REFERENCES

Albert, D.Z. (1992). *Quantum mechanics and experience*. Cambridge, Mass.: Harvard University Press.

Allen, K. (2010a). 'Locating the unique hues', Rivista di Estetica 43: 13-28.

Allen, K. (2010b). 'In defence of daylight', Pacific Philosophical Quarterly 91: 1-18.

Audi, P. (2012). 'A clarification and defense of the notion of grounding', in Correia, F. & Schnieder, B. (eds.), *Metaphysical Grounding: Understanding the Structure of Reality*. Cambridge: Cambridge University Press.

Austin, J.L. (1962). Sense and Sensibilia. Oxford: Oxford University Press.

Barker, S. & Jago, M. (2014). 'Monism and material constitution', *Pacific Philosophical Quarterly* 95: 189-204.

Batty, C. (2010). 'Olfactory experience II: objects and properties', *Philosophy Compass* 5/12: 1147-1156.

Billock, V.A. & Tsou, B.H. (2010). 'Seeing forbidden colors', *Scientific American* 302(2): 72-77.

Block, N. (2003). 'Mental paint', in Hahn, M. & Ramberg, B. (eds.), *Essays on the Philosophy of Tyler Burge*. Cambridge, Mass.: The MIT Press.

Block, N. (1999). 'Sexism, ageism, racism, and the nature of consciousness', *Philosophical Topics* 26(1): 39-70.

Block, N. (1990). 'Inverted Earth', Philosophical Perspectives 4: 53-79.

Brewer, B. (2011). Perception and Its Objects. Oxford: Oxford University Press.

Brewer, B. (2007). 'Perception and its objects', Philosophical Studies 132: 87-97.

Brewer, B. (2006). 'Perception and content', *European Journal of Philosophy* 14(2): 165-181.

Byrne, A. & Hilbert, D. (2007). 'Colour primitivism', Erkenntnis 66: 73-105.

Byrne, A. & Logue, H. (2008). 'Either/or', in Haddock, A. & Macpherson, F., *Disjunctivism: Perception, Action, Knowledge*. Oxford: Oxford University Press.

Byrne, A. & Tye, M. (2006). 'Qualia ain't in the head', Noûs 40(2): 241-255.

Campbell, J. (2002a). Reference and Consciousness. Oxford: Clarendon Press.

Campbell, J. (2002b). 'Berkeley's puzzle', in Gendler, T.S. & Hawthorne, J. (eds.), *Conceivability and Possibility*. Cambridge, Mass.: The MIT Press.

Chalmers, D. (2006). 'Perception and the fall from Eden', in Gendler, T.S. & Hawthorne, J. (eds.), *Perceptual Experience*. Oxford: Oxford University Press.

Chalmers, D. (2002). 'Consciousness and its place in nature', in Chalmers, D. (ed.), *Philosophy of Mind: Classic and Contemporary Readings*. Oxford: Oxford University Press.

Child, W. (1994). 'Vision and causation: reply to Hyman', *The Philosophical Quarterly* 44: 361-369.

Child, W. (1992). 'Vision and experience: the causal theory and the disjunctive conception', *The Philosophical Quarterly* 42(168): 297-316.

Cohen, J. (2010). 'It's not easy being green: Hardin and color relationalism', in Cohen, J. & Matthen, M. (eds.), *Color Ontology and Color Science*. Cambridge, Mass.: The MIT Press.

Cohen, J. (2007). 'A relationalist's guide to error about color perception', *Noûs* 41(2): 335-353.

Cohen, J. (2004). 'Color properties and color ascriptions: a relationalist manifesto', *The Philosophical Review* 113(4): 451-506.

Cohen, J., Hardin, C.L. & McLaughlin, B.P. (2006). 'True colours', *Analysis* 66(4): 335-340.

Conway, C.M. & Stoughton, B.R. (2008). 'Neural basis for unique hues', *Current Biology* 18(16): R698-R699.

Crane, H.D. & Piantanida, T.P. (1983). 'On seeing reddish green and yellowish blue', *Science* 221: 1078-1080.

Daly, C. (2012). 'Scepticism about grounding', in Correia, F. & Schnieder, B. (eds.), *Metaphysical Grounding: Understanding the Structure of Reality*. Cambridge: Cambridge University Press.

Davidson, D. (2000). 'A coherence theory of truth and knowledge', in Sosa, E. & Kim, J. (eds.) *Epistemology: An Anthology*. Oxford: Blackwell Publishing Ltd. [check original (Lepore)]

De Valois, R.L. & De Valois, K.K. (1997). 'Neural coding of color', in Byrne, A. & Hilbert, D.R. (eds.), *Readings on Color Volume 2: The Science of Color*. Cambridge, Mass.: The MIT Press.

Delk, J.L. & Fillenbaum, S. (1965). 'Differences in perceived color as a function of characteristic color', *The American Journal of Psychology* 78(2): 290-293.

Dowden, B. (2013). 'Time', *Internet Encyclopedia of Philosophy*. URL = http://www.iep.utm.edu/time/

Dowe, P. (2009). 'Causal process theories', in Beebee, H., Hitchcock, C. & Menzies, P. (eds.), *The Oxford Handbook of Causation*. Oxford: Oxford University Press.

Dowe, P. (1992). 'Wesley Salmon's process theory of causality and the conserved quantity theory', *Philosophy of Science* 59(2): 195-216.

Dretske, F. (2003). 'Experience as representation', Philosophical Issues 13: 67-82.

Dretske, F. (2000). 'Simple seeing', in *Perception, Knowledge and Belief: Selected Essays*. Cambridge: Cambridge University Press.

Dretske, F. (1995). 'Phenomenal externalism or if meanings ain't in the head, where are qualia?', *Philosophical Issues* 7: 143-158.

Dunn, J. (2008). 'The obscure act of perception', Philosophical Studies 139: 367-393.

Ekroll, V. & Faul, F. (2012). 'Basic characteristics of simultaneous color contrast revisited', *Psychological Science* 23(10): 1246-1255.

Esfeld, M. (1999). 'Physicalism and ontological holism', Metaphilosophy 30(4): 319-337.

Fair, D. (1979). 'Causation and the flow of energy', Erkenntnis 14(3): 219-250.

ffytche, D.H., Howard, R.J., Brammer, M.J., David, A., Woodruff, P. & Williams, S. (1998). 'The anatomy of conscious vision: an fMRI study of visual hallucinations', *Nature Neuroscience* 1(8): 738-742.

Fish, W. (2013). 'Perception, hallucination, and illusion: reply to my critics', *Philosophical Studies* 163: 57-66.

Fish, W. (2010). *Philosophy of Perception: A Contemporary Introduction*. Abingdon: Routledge.

Fish, W. (2009). Perception, Hallucination, and Illusion. Oxford: Oxford University Press.

Fodor, J. (1989). Psychosemantics. Cambridge, Mass.: The MIT Press.

Gibson, E.J. (1963). 'Perceptual learning', Annual Review of Psychology 14: 29-56.

Gibson, J.J. & Gibson, E.J. (1955). 'Perceptual learning: differentiation or enrichment?', *Psychological Review* 52(1): 32-41.

Glüer, K. (2007). 'Colors without circles?', Erkenntnis 66: 107-131.

Grice, H.P. (1961). 'The causal theory of perception', *Proceedings of the Aristotelian Society, Supplementary Volumes*, 35: 121-152.

Grimes, T.R. (1991). 'Supervenience, determination, and dependency', *Philosophical Studies* 62: 81-92.

Hamlyn, D.W. (1983). *Perception, Learning and the Self: Essays in the Philosophy of Psychology*. London: Routledge & Kegan Paul.

Hardin, C.L. (1997). 'Reinverting the spectrum', in Byrne, A. & Hilbert, D.R. (eds.), *Readings on Color Volume 1: The Philosophy of Color*. Cambridge, Mass.: The MIT Press.

Hardin, C.L. (1988). *Color for Philosophers: Unweaving the Rainbow*. Indianapolis: Hackett Publishing Company.

Harman, G. (1996). 'Explaining objective color in terms of subjective reactions', *Philosophical Issues* 7: 1-17.

Harman, G. (1990). 'The intrinsic quality of experience', *Philosophical Perspectives* 4: 31-52.

Hansen, T., Olkkonen, M., Walter, S. & Gegenfurtner, K.R. (2006). 'Memory modulates color appearance', *Nature Neuroscience* 9(11): 1367-1368.

Hilbert, D. (2004). 'Hallucination, sense-data and direct realism', *Philosophical Studies* 120: 185-191.

Howard, J.D., Plailly, J., Grueschow, M., Haynes, J-D. & Gottfried, J.A. (2009). 'Odor quality coding and categorization in human posterior piriform cortex', *Nature Neuroscience* 12: 932-938.

Hume, D. (2007). *An Enquiry Concerning Human Understanding* (Peter Millican, ed.). Oxford: Oxford University Press.

Hyman, J. (1994). 'Reply to Vision', The Philosophical Quarterly 44: 369-376.

Hyman, J. (1992). 'The causal theory of perception', *The Philosophical Quarterly* 42: 277-296.

Jameson, D. & Hurvich, L.M. (1959). 'Perceived color and its dependence on focal, surrounding, and preceding stimulus variables', *Journal of the Optical Society of America* 49(9): 890-898.

Johnston, M. (2004). 'The obscure object of hallucination', *Philosophical Studies* 120: 113-183.

Johnston, M. (1992). 'How to speak of the colors', Philosophical Studies 68(3): 221-263.

Kalderon, M. (2007). 'Colour pluralism', Philosophical Review 116(4): 563-601.

Kay, P. & McDaniel, C.K. (1997). 'Linguistic significance of meanings of basic color terms', in Byrne, A. & Hilbert, D.R. (eds.), *Readings on Color Volume 2: The Science of Color*. Cambridge, Mass.: The MIT Press.

Keller, A., Zhuang, H., Chi, Q., Vosshall, L.B. & Matsunami, H. (2007). 'Genetic variation in a human odorant receptor alters odour perception', *Nature* 449(7161): 468-72.

Kennedy, M. (2009). 'Heirs of nothing: The implications of transparency', *Philosophy and Phenomenological Research* 79(3): 574-604.

Kitayama, S. & Cohen, D. (2007). *Handbook of Cultural Psychology*. New York: The Guilford Press.

Kuehni, R.G. (2004). 'Variability in unique hue selection: a surprising phenomenon', *Color Research and Application* 29(2): 158-162.

Levin, D.T. & Banaji, M.R. (2006). 'Distortions in the perceived lightness of faces: the role of race categories', *Journal of Experimental Psychology: General* 135(4): 501-512.

Lewis, D. (1986). On the Plurality of Worlds. Oxford: Blackwell.

Locke, J. (1993). *An Essay Concerning Human Understanding*, ed. J. Yolton. London: Everyman.

Lotto, R.B. & Purves, D. (2000). 'An empirical explanation of color contrast', *Proceedings* of the National Academy of Sciences of the United States of America 97(23): 12834-12839.

Lowe, E.J. (2008). 'Against disjunctivism', in Haddock, A. & Macpherson, F., *Disjunctivism: Perception, Action, Knowledge*. Oxford: Oxford University Press.

Lycan, W. (2001). 'The case for phenomenal externalism', Noûs 35(s15): 17-35.

Macpherson, F. (2012). 'Cognitive penetration of colour experience: rethinking the issue in light of an indirect mechanism'. *Philosophy and Phenomenological Research* 84(1): 24-62.

Manzotti, R. (2011). 'The spread mind: seven steps to situated consciousness', *Journal of Cosmology* 14: 4526-4535.

Manzotti, R. (2008). 'A process-oriented view of qualia', in Wright, E. (ed.), *The Case for Qualia*. Cambridge, Mass.: The MIT Press.

Manzotti, R. (2006a). 'A process-oriented view of conscious perception', *Journal of Consciousness Studies* 13(6): 7-41.

Manzotti, R. (2006b). 'Consciousness and existence as a process', *Mind & Matter* 4(1): 7-43.

Markosian, N. (2014). 'Time', *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition), Edward N. Zalta (ed.), URL = ">http://plato.stanford.edu/archives/spr2014/entries/time/.

Martin, M.G.F. (2013). '*Shibboleth:* some comments on William Fish's *Perception, Hallucination & Illusion'*, *Philosophical Studies* 163: 37-48.

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: 37-89.

Maudlin, T. (2011). *Quantum non-locality and relativity (third edition)*. Chichester, West Sussex: Wiley-Blackwell.

Maund, B. (2012). 'Color', *The Stanford Encyclopedia of Philosophy* (Winter 2012 Edition), Edward N. Zalta (ed.), URL = ">http://plato.stanford.edu/archives/win2012/entries/color/.

McDermid, D.J. (2001). 'What is direct perceptual knowledge? A fivefold confusion', *Grazer Philosophische Studien* 62: 1-16.

McDowell, J. (2009). *Having the World in View: Essays on Kant, Hegel, and Sellars*. Cambridge, Mass.: Harvard University Press.

McDowell, J. (1994). Mind and World. Cambridge, Mass.: Harvard University Press.

McTaggart, J.M.E. (1908). 'The unreality of time', Mind 17: 457-473.

Miller, A. (2010). 'Realism', in Zalta, E.N. (ed.), *The Stanford Encyclopedia of Philosophy* (*Spring 2012 Edition*). URL = ">http://plato.stanford.edu/archives/spr2012/entries/realism/.

Miller, K. (2013). 'Presentism, eternalism, and the growing block', in Dyke, H. & Bardon, A. (eds.), *A Companion to the Philosophy of Time*. Oxford: Wiley-Blackwell.

O'Callaghan, C. (2014). 'Auditory perception', *The Stanford Encyclopedia of Philosophy* (Summer 2014 Edition), Edward N. Zalta (ed.), forthcoming URL = http://plato.stanford.edu/archives/sum2014/entries/perception-auditory/>.

Pautz, A. (2014). 'The real trouble with phenomenal externalism: 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.

Pautz, A. (2013). 'Do the benefits of naïve realism outweigh the costs? Comments on Fish, *Perception, Hallucination and Illusion*', *Philosophical Studies* 163: 25-36.

Pautz, A. (2011). 'Can disjunctivists explain our access to the sensible world?, *Philosophical Issues* 21: 384-432.

Pears, D.F. (1976). 'The causal conditions of perception', Synthese 33: 25-40.

Power, S.E. (2010). 'Perceiving external things and the time-lag argument', *European Journal of Philosophy* 21(1): 94-117.

Ramsey, W. (2013). 'Eliminative materialism', *The Stanford Encyclopedia of Philosophy* (Summer 2013 Edition), Edward N. Zalta (ed.), URL = http://plato.stanford.edu/archives/sum2013/entries/materialism-eliminative/>.

Rea, M. (2005). 'Four-dimensionalism', in Loux, M.J & Zimmerman, D.W. (eds.), *The Oxford Handbook of Metaphysics*. Oxford: Oxford University Press.

Robinson, H. (1994). Perception. London: Routledge.

Salmon, W.C. (1994). 'Causality without counterfactuals', *Philosophy of Science* 61(2): 297-312.

Salmon, W.C. (1984). *Scientific Explanation and the Causal Structure of the World*. Princeton, N.J.: Princeton University Press.

Schaffer, J. (2010). 'Monism: the priority of the whole', *Philosophical Review* 119(1): 31-76.

Schefrin, B.E. & Werner, J.S. (1990). 'Loci of spectral unique hues throughout the lifespan', *Journal of the Optical Society of America A* 7(2): 305-311.

Schellenberg, S. (2013). 'Externalism and the gappy content of hallucination', in Macpherson, F. & Platchias, D. (eds.), *Hallucination: Philosophy and Psychology*. Cambridge, Mass.: The MIT Press.

Sellars, W. (2000). 'Empiricism and the philosophy of mind', reprinted in deVries, W.A. & Triplett, T., *Knowledge, Mind, and the Given: Reading Wilfrid Sellars's "Empiricism and the Philosophy of Mind"*. Indianapolis: Hackett Publishing Company, Inc.

Shoemaker, S. (1991). 'Qualia and consciousness', Mind 100(4): 507-524.

Sider, T. (2001). *Four-Dimensionalism: An Ontology of Persistence and Time*. Oxford: Oxford University Press.

Siegel, S. (2008). 'The epistemic conception of hallucination', in Haddock, A. & Macpherson, F., *Disjunctivism: Perception, Action, Knowledge*. Oxford: Oxford University Press.

Skow, B. (2009). 'Relativity and the moving spotlight', *The Journal of Philosophy* 106(12): 666-678.

Snowdon, P. (1981). 'Perception, vision, and causation', *Proceedings of the Aristotelian Society*, New Series, 81: 175-192.

Snowdon, P. (1990). 'The objects of perceptual experience', *Proceedings of the Aristotelian Society, Supplementary Volumes* 64: 121-150.

Sollberger, M. (2008). 'Naïve realism and the problem of causation', *Disputatio* 3(25): 1-19.

Strawson, P.F. (2011). 'Perception and its objects', in *Philosophical Writings*. Oxford: Oxford University Press.

Thoen, H.H., How, M.J., Chiou, T-H, Marshall, J. (2014). 'A different form of colour vision in mantis shrimp'. *Science* 343: 411-413.

Thompson, E. (1995). *Colour Vision: A Study in Cognitive Science and the Philosophy of Perception*. Abingdon, Oxon.: Routledge.

Travis, C. (2007). 'Reason's reach', European Journal of Philosophy 15(2): 225-248.

Travis, C. (2004). 'The silence of the senses', Mind 113(449): 57-94.

Tye, M. (2009). *Consciousness Revisited: Materialism Without Phenomenal Concepts*. Cambridge, Mass.: The MIT Press.

Tye, M. (2006). 'The puzzle of true blue', Analysis 66(3): 173-178.

Valberg, J.J. (1992). 'The puzzle of experience' in Crane, T. (ed.), *The Contents of Experience: Essays on Perception*. Cambridge: Cambridge University Press.

Webster, M.A., Miyahara, E., Malkoc, G. & Raker, V.E. (2000). 'Variations in normal color vision II: unique hues', *Journal of the Optical Society of America A* 17(9): 1545-1555.

Weir, A. (2004). 'An ultra-realist theory of perception', *International Journal of Philosophical Studies* 12(2): 105-128.