

An Action-Space Theory of Conscious Vision

Dave Ward

**PhD in Philosophy
The University of Edinburgh
2008**

I have read and understood the University of Edinburgh's regulations on plagiarism. All material contained herein is my own work, except where indicated otherwise by way of quotations, references and acknowledgements.

Acknowledgements

Thanks first and foremost to my supervisors, Andy Clark and Pepa Toribio for helpful comments, suggestions, advice and more at every stage of the preparation of this thesis. I have also benefited greatly from comments by, and discussions with, Tim Bayne, Julian Kiverstein, Alva Noë, Matt Nudds and Tom Roberts. Thanks also to the AHRC for providing the funding that has allowed me to carry out the research for this thesis. Finally, I owe a special debt of gratitude to Susan Hurley, with whom I discussed much of the material in this thesis during a few weeks I spent at the University of Bristol, shortly before her death. She was and remains an inspiration to me both professionally and personally, and I count my opportunity to have worked with her as one of the great privileges of my life.

Abstract

I argue that conscious visual experience consists in a direct and non-inferential grasp of the way one's current perceptual contact with the environment poises one to pursue various intentional plans, goals and projects. I show that such a view of visual consciousness is supported by current work in cognitive neuroscience, affords a compelling account of colour perception, and suggests a way to bridge the 'explanatory gap' between consciousness and the language of the natural sciences. In chapter 1, I examine the reasoning that leads to the appearance of an explanatory gap between the phenomenal and the physical in more detail, and set out the constraints on a solution that our discussion of the problem has imposed. I then sketch the two rival takes on the relationship between perception and action mentioned above – adjudicating between these two theories (and finding in favour of the action-space view) is the task of the next two chapters, and is a recurring theme throughout. Chapter 2 moves on to discuss some recent work in the neuropsychology of vision and what it might suggest about the functional role of conscious vision, and the first half of chapter 3 considers two puzzle cases concerning colour perception. Each of these discussions turns out to constitute a source of support for the action-space view that visual perception consists in a grasp of the practical consequences of sensation, and the second half of chapter 3 sets out this view and responds to an initial range of questions and objections it might face. Chapter 4 illustrates our view via a discussion of colour perception, and chapter 5 discusses the type of grasp of practical consequences that is necessary for perceptual sensitivity to issue in conscious experience. By chapter 6, we are in a position to see how the action-space approach can help close the explanatory gap for phenomenal consciousness, and our final chapter sets out how I think this should be done. I conclude with a brief discussion of further questions and prospects for the action-space approach.

Introduction

This is a dissertation about conscious experience. More specifically, it is about two sorts of questions about experience: firstly, is it possible to give a reductive explanation of conscious experience – to explain what it is and why it occurs in terms of anything less than conscious experience itself? Secondly, what role should an appeal to involved, intelligent action play in an account of conscious experience? It seems striking that whenever we find a locus of suitably intelligent action geared towards successfully inhabiting and navigating an environment, we are inclined to suppose that we have also found a seat of conscious experience – how might such intuitions be grounded, and how should they inform an account of experience?

Much ink has been spilt over each of these questions in recent years. Our first question (which we'll consider in detail very shortly) lies at the heart of the debate concerning the apparent 'explanatory gap' between conscious experience and the sorts of neural, psychological, functional and physical properties in terms of which we might try to ground an explanation of that experience. As we shall see in chapter 1, there are various ways to bring out the worry that no amount of information about the latter sorts of properties can ever add up to complete information about the qualitative character of a conscious experience – the way that experience feels to the person who undergoes it – but here is one. Einstein is alleged to have said that science cannot give us the taste of chicken soup. Presumably he meant that although the physical science of soup might allow us to deduce many facts about its chemical composition, its colour, its viscosity, its temperature, and innumerable other properties of the soup, it will not allow us to deduce important facts about the particular way the soup tastes to its eater. We might, given our accumulated scientific knowledge about the various ways in which people are disposed to react to soups with certain physical properties, be able to predict that the eater is likely to describe the soup as tart, or tepid, or thin, but this, it seems does not amount to complete information about *how eating the soup actually feels* to its consumer. We might conclude, then, that our experiences of eating soup have certain sorts of properties that science cannot tell us about – the properties that specify what it's like to undergo that experience. We can call these properties the experience's *qualitative character*, or its *phenomenal properties*, or *qualia*. Such properties seem to be stubborn obstacles to an explanation of experience in physical terms – they seem to be essentially ineffable, private and subjective, in contrast to the public and objective language in terms of which we wish to explain them. We have just suggested, in fact, that these properties do not seem to be entailed by any combination of the experience's non-phenomenal properties. An apparently innocent line of reasoning seems to

have led to the conclusion that experiences have kernels of pure phenomenal properties that cannot be accounted for in terms of the non-phenomenal properties of experience. Part of the aim of this thesis is to examine whether and how such a view of experience can be resisted.

Our second question concerned the nature of the appeal that an account of conscious experience should make to skilled behaviour. In contemporary philosophy of mind, there is a strong current that emphasises the embodied, environmentally extended and active nature of minded creatures. There is a growing camp of theorists from philosophy, the cognitive sciences and related disciplines all concerned to emphasize that we should view the mind as (to borrow from the title of a recent conference) embodied, embedded, enacted and extended. A multiplicity of different empirical and conceptual work tends to be grouped together in support of this vision, and I won't attempt the difficult tasks of specifying lines of thought that unite such work, or teasing apart the many different strands of thought and research that divide it. I'll be concerned solely with investigating the proper role of an appeal to intelligent action in an account of conscious visual experience. This is a question that has been pursued by a number of different theorists in the above camp. One thing that many such theorists have in common is a suspicion that a certain 'traditional' view of the relations between perception, cognition and action is in need of revision. According to this traditional view, (dubbed "the Classical Sandwich" by Susan Hurley) perception takes in raw data from the environment, which is then operated on by some central cognitive faculties, that then issue commands for actions in light of those cognitive operations. However, there is a possibility that due stress on the embodied, embedded and active nature of mind might suggest that relations between perception, cognition and action are more intricate and interdependent, and less serial and segregated, than the traditional picture would have us believe.

One way of fleshing out this possibility is suggested by work on "sensorimotor" theories of perception. According to such theories, perception consists in a practical grasp of the sensory consequences of movement. Perceiving the shape of a visually presented cup, for example, consists in understanding the various ways in which possible movements are related to possible sensory impressions. Perceiving a vertical line consists partly in understanding that nodding one's head up and down will leave visual sensations of the line unchanged, and understanding the predictable and regular way in which moving one's head from side to side will cause sensations to change. On such a view, perception and action are intimately related, rather than the occupants of two extremes in a perception – action – cognition progression. Part of the aim of this thesis is to argue that this view of

the intimate relation between perception and action is mistaken. Whilst the sensorimotor theorist understands visual perception in terms of a grasp of the sensory consequences of action, I will argue that we should understand it in terms of a grasp of the practical consequences of sensation. According to the “action-space” account I’ll develop in what follows, visual perceptual experience consists in a certain kind of openness to the possibilities for action afforded by the perceiver’s current environment. The chief aim in this thesis is to elaborate on this claim, and marshal both empirical and philosophical reasons for accepting it.

Importantly, I think that the two questions we’ve just briefly considered should be addressed together. I will suggest that an appeal to the essentially active nature of perception can show us how to deal with worries about the apparent inexplicability of visual experience in non-experiential terms. And considering how alternative ways of emphasising the role of action in perception interface with questions about the explicability of conscious experience can provide us with reasons to prefer one sort of emphasis over another. The action-space account of visual experience I develop suggests how we might close the explanatory gap, and the fact that it does so constitutes one strong reason (among others) to think that it constitutes the correct view of the relationship between perceptual experience and skilled action.

Here is how we’ll proceed. In chapter 1, I examine the reasoning that leads to the appearance of an explanatory gap between the phenomenal and the physical in more detail, and set out the constraints on a solution that our discussion of the problem has imposed. I then sketch the two rival takes on the relationship between perception and action mentioned above – adjudicating between these two theories (and finding in favour of the action-space view) is the task of the next two chapters, and is a recurring theme throughout. Chapter 2 moves on to discuss some recent work in the neuropsychology of vision and what it might suggest about the functional role of conscious vision, and the first half of chapter 3 considers two puzzle cases concerning colour perception. Each of these discussions turns out to constitute a source of support for the action-space view that visual perception consists in a grasp of the practical consequences of sensation, and the second half of chapter 3 sets out this view and responds to an initial range of questions and objections it might face. Chapter 4 illustrates our view via a discussion of colour perception, and chapter 5 discusses the type of grasp of practical consequences that is necessary for perceptual sensitivity to issue in conscious experience. By chapter 6, we are in a position to see how the action-space approach can help close the explanatory gap for phenomenal

consciousness, and our final chapter sets out how I think this should be done. I conclude with a brief discussion of further questions and prospects for the action-space approach.

Let's begin...

Chapter 1: Deflating Experiences

Looking at the sky outside the window by my desk, I have an experience of its colour. Focussing on this experience, it seems natural to say that the sky's colour qualitatively looks a certain way to me, or in plainer language, that my experience of the sky's colour *feels* a certain way to me. And, generalising from this observation, it seems natural to think that an experience's feeling one way rather than another to its subject is an *essential property* of that experience. One way of individuating experiences is by reference to such properties – the property of what it feels like to have that experience. Philosophers have labelled such properties of experience *phenomenal properties*, or *qualia*, and the sum of these properties for a particular experience the *phenomenology* of that experience. Someone who has experiences with such properties enjoys, they say, *phenomenal consciousness*.

But whilst it seems natural to say that there are both experiences and such properties of experiences, this view of matters appears to conflict with an equally natural view of the world and our place in it. According to what we can call a 'minimal naturalism', all the world's properties are either properties posited by natural sciences, or explicable in terms of those properties¹. The explanatory and predictive power of the natural sciences and the lack of counterexamples to such a minimal naturalism drawn from outside the realm of experience give us good reason to adopt such a worldview. But when we attempt to find a place for experience and its phenomenal properties in such a view, we appear to run into difficulties.

1.1 The Phenomenal and the Physical

A stark way to bring out these difficulties is suggested by Chalmers (1996, 2003), who argues as follows:

P1: Physical accounts explain at most structure and function.

P2: Explaining structure and function does not suffice to explain consciousness.

C: No physical account can explain consciousness.

(Chalmers (2003), p.105)

¹ This rough definition of naturalism suffices for my purposes here. For a thorough discussion, see MacDonald (1992).

It seems plausible, as P1 claims, that physical facts² can, at most, tell us about structural and functional organisation. All we can deduce from information about physical properties is information about structural and functional properties. But reflecting on the phenomenal properties of my experience of the sky's colour yields the intuition that these are not structural or functional properties. A specification of the structural and functional properties of my experience, such as its neurophysiological basis and its relations to various discriminatory and motivational states, does not seem to yield information that suffices to determine its phenomenal properties – how the experience feels to me (P2). Accepting this premise is what gives the argument its anti-physicalist force. It relies upon the intuition that simply reflecting on a given experience allows us to see that no amount of structural or functional information can add up to information that determines the experience's phenomenal properties. If we credit this intuition, Chalmers' conclusion follows: phenomenal consciousness cannot be explained in physical terms. Here are three ways of making the same point.

The Zombie Argument: Chalmers (1996 p.94-9) invites his readers to consider a creature physically identical to them, but lacking phenomenal consciousness. When my zombie twin contemplates the colour of the sky, he processes optical information identically to me, and is disposed to judge, discriminate and act in an identical manner to me as a result of this processing, but lacks any phenomenal experience of the sky's colour. If we can conceive of our zombie twins without apparent incoherence or contradiction, this shows that we cannot see how the physical facts about us serve to fix or explain the phenomenal facts about us. If we could see how they did so, then we could see that the physical identity that obtains between us and our zombie twins precludes there being any phenomenal difference between us – and thus that the possibility of a creature's sharing our physical properties whilst lacking our phenomenal properties³ is incoherent. The apparent conceivability of zombies shows that we lack a grasp of how our physical properties fix or explain the phenomenal properties of our experiences.

² References to physical facts and properties should be understood to encompass all facts and properties compatible with a minimal naturalism (e.g. to encompass functional and intentional facts and properties). This is to accommodate skepticism about whether all truths of the natural sciences are reducible to truths of physics. See Fodor (1974) and Laurence and Margolis (2003) for reasons for such skepticism.

³ Strictly speaking, phenomenal properties are properties of experiences rather than of persons. Throughout, references to "S's phenomenal properties" should be read as abbreviating "The phenomenal properties of the experience S is currently having."

The Inverted Spectrum Argument⁴: Alternatively, we can ask whether we can conceive of a physical duplicate of ourselves whose experiences have phenomenal properties that differ from those of my experiences. In virtue of our physical identity, my duplicate and I will act, discriminate and judge in an identical manner, including exhibiting identical colour naming behaviour, and judging similarity relations between colours to be the same. In light of this, a simple switching of the phenomenal properties of yellow and blue experiences alone would not result in the indistinguishable inversion between my twin and I that we require – our judgements, dispositions and discriminatory capacities concerning the relations between yellow, blue, and the intervening colours would also need to be systematically inverted. A widespread intuition is that such an inversion of phenomenal properties between physical or functional duplicates is more readily conceivable than the zombie scenario above. But, as with the zombie argument, the lack of apparent incoherence in such a situation shows that we lack a grasp of how our physical properties fix our phenomenal properties. Our accepting the conceivability of zombies requires us to acknowledge that we do not grasp how our physical properties ensure that we have experiences with phenomenal properties. Accepting the conceivability of inverted spectra requires us to acknowledge that we lack a grasp of how physical properties ensure that the phenomenal properties of our experiences are one way rather than another.

The Knowledge Argument: Lastly, the knowledge argument (Jackson (1982)) invites us to consider Mary, a brilliant colour scientist who has spent her life incarcerated in a black and white room. Over the course of her life, Mary has acquired a comprehensive knowledge of all the physical properties of the world, via reading black and white textbooks and watching black and white television. In particular, she has a complete knowledge of the physical properties of coloured objects and of light, and of the precise way such information is processed by human colour-vision. But the nature of her surroundings has ensured that Mary has experienced no colours other than black, white and intervening shades of grey. One day, she is released into the coloured world, and has her first colour-experience whilst looking at a red rose. Our intuitions say that this experience will constitute a discovery for Mary – her learning a new fact about what it is like to experience a red object. But all the relevant physical facts about roses, light and the mechanics of her colour-vision were already in her possession. So, we are invited to conclude, the new fact Mary learns about experience is not a physical fact. Complete knowledge of physical facts (and complete ability to make relevant deductions from those

⁴ Shoemaker (1982) is the locus classicus for contemporary discussions of inverted spectra. See also Palmer (1999) for an interesting list of possible (and impossible) inversions and their relative likelihoods.

facts) does not entail knowledge of all the facts about what it's like to have a particular experience. Our intuition that Mary will learn something new upon leaving her room shows that we lack a grasp of how knowledge of physical properties could be translated into knowledge of phenomenal properties.

Each of these arguments tests the connections between our concepts of phenomenal consciousness and those of the physical, structural and functional. If we accept the intuition behind the knowledge argument, or the conceivability of zombies or inverted spectra, it seems that we must conclude that there is an 'explanatory gap' between our concept of consciousness and the physical concepts with which we wish to explain it. Following Carruthers (2004), we might say that there is no 'conceptual mesh' between the explanandum of consciousness and any candidate physical explanans. Were there such a mesh, we could see how Mary could move from her complete knowledge of physical facts to the knowledge of phenomenal facts she intuitively lacks, and how the physical identity between myself and my putatively inverted or zombie twin rules out there being any phenomenological difference between us. But, proponents of the explanatory gap argue, it is just this sort of conceptual mesh that reductive explanation aims to provide. If a physicist claims to have reductively explained temperature in terms of mean molecular kinetic energy, but also that she can coherently conceive of holding facts about mean molecular kinetic energy fixed whilst facts about temperature vary, or that she might be in possession of complete information about something's mean molecular kinetic energy (plus a suitably keen mind) whilst unable to deduce its temperature, we should conclude either that she has misunderstood her own explanation, or that her explanation has failed⁵. A good explanation of temperature in terms of kinetic energy would let us see how to move from the facts about kinetic energy to the facts in the target domain – temperature. Being able to conceive of the capability of one set of facts varying independently of the other shows that we do not understand how to move from the domain in terms of which we wish to explain the target phenomenon and the target phenomenon in the required way. For this

⁵ This might not be so if our scientist was a sort of Humean, believing that facts about temperature and facts about mean molecular kinetic energy were merely constantly correlated. But this doesn't seem like a desirable form for a reductive explanation of consciousness. It rules out arguing that phenomenal facts are identical to structural/functional facts, or that they can be deduced from those facts independently of knowledge of brute bridging principles. The undesirability of an appeal to such principles is addressed in the next section's discussion of dualism and type-B materialism. Relatedly, these discussions provide support for my endorsement of the 'conceptual mesh' requirement.

reason, it seems that there is an explanatory gap between the phenomenal and the physical⁶.

1.2 Mapping the Gap

How should we respond to the arguments for the existence of an explanatory gap? One response is to reject the intuitions required for the arguments of the previous section to work – the conceptual mesh required for reductive explanation does indeed obtain between the phenomenal and physical domains. As a result, zombies and inverted spectra are not really conceivable, and Mary learns nothing new upon her release. Upon reflection, we can see that there is nothing in our concept of phenomenal consciousness that prohibits its explanation in physical terms. Chalmers (2003) calls such a position ‘**type-A materialism**’, since the claim that there is the required conceptual mesh between the phenomenal and the physical amounts to a claim of *a priori* entailment between the two domains. Philosophers occupying this camp will include analytic functionalists such as (one interpretation of) Dennett (1991), who argue that our concept of consciousness only picks out various functional and intentional activity⁷, and eliminative materialists such as Rorty (1965), who argue that our concept of consciousness picks out nothing at all.

Sticking with Chalmers’ taxonomy, ‘**type-B materialism**’ claims that phenomenal properties are identical to physical properties, but that this identity is not knowable *a priori*. Rather, the links which obtain between the physical and phenomenal are to be discovered *a posteriori*, by empirical investigation rather than conceptual analysis. The relationship between consciousness and its physical basis should be understood on the model of other empirically discovered identities, such as water’s identity with H₂O⁸. Such a position can acknowledge that we share the intuitions relied upon by the above arguments whilst denying that this is a barrier to a physicalist account of consciousness – zombies and inverted spectra seem conceivable and Mary seems partially ignorant because we lack an *a priori* grasp of the relevant identities. But the identities between phenomenal and physical states are there to be discovered empirically nonetheless. Type-B

⁶ For a defence of the conception of reductive explanation at work here, see Chalmers and Jackson (2001). For criticism of this conception, see Laurence and Margolis (2003).

⁷ Though Dennett’s overall position is difficult to fit into any philosophical taxonomy, I take the spirit of his view to be consistent with the position I eventually defend rather than the analytic functionalism or eliminativism he is sometimes taken to endorse.

⁸ Alternatively, type-B materialists can hold that the phenomenal is *a posteriori* necessitated by, rather than identical to, the physical. The rest of my discussion is framed in terms of identities, but all arguments are neutral between the two options.

materialism thus allows the attractive combination of conceptual dualism with respect to the phenomenal and the physical, with ontological physicalism. The dualist component accommodates our intuitions about our epistemic situation with regard to consciousness whilst the physicalist component accommodates our naturalist tendencies. Psycho-functionalists and empirical functionalists such as Putnam (1960/75) and Lycan (1996) will occupy this camp, as well as the representationalist and phenomenal concept strategies of Tye (1995) and Carruthers (2007), to be discussed in chapter 6.

Type-C materialists hold that the explanatory gap can be closed in principle, but that we currently lack the theoretical or cognitive resources necessary to see how this is possible. Some theorists hold that this lack is a necessary and permanent consequence of our cognitive limitations (McGinn (1991)), whilst others suggest that conceptual (Nagel (1974)) or empirical (Churchland (1997)) advances could resolve matters.

Lastly, **dualists** of various types hold that our difficulty in seeing how consciousness might be explained in physical terms is due to the fact that consciousness is a feature of the universe that is incompatible with our current conception of physicalism. As a result, the physicalist ontology should be either abandoned (Foster (1989)) or expanded to encompass fundamental phenomenal properties (e.g. Chalmers (1996)).

It seems to me that none of the above theories are promising candidates for a satisfactory explanation of phenomenal consciousness. The position that fails most clearly is type-C materialism, which asserts that such an explanation is impossible. The position of optimistic type-C materialists such as Nagel and Churchland collapses into type-A or B materialism (each of which are criticised below), depending on the type of explanation available once the relevant conceptual or scientific advances have been made.

Dualism posits a new fundamental feature of the universe in order to account for phenomenal properties. As well as being ontologically extravagant, this gives up on the prospect of demonstrating a conceptual mesh between the phenomenal and the physical, a necessary criterion for reductive explanation as I have construed it here. One option for the dualist is to claim that the fundamental phenomenal properties they posit are necessarily immune to our investigation, clearly ruling out the possibility of demonstrating a mesh between those properties and familiar physical properties. Alternatively, dualists can posit fundamental laws linking phenomenal and physical properties (Chalmers (2004)). But whilst such laws might be empirically investigated by discovering correlations between phenomenal and physical properties, we could never arrive at a satisfactory

understanding of why one set of laws rather than another obtains. In order to do so, we would need to understand how things being a certain way physically necessitates their being a certain way phenomenally (or vice versa), otherwise the laws linking the two domains will appear contingent. But in doing so, we would have obtained the desired conceptual mesh between the phenomenal and the physical, and our position would be a type-A materialism – there would be no need to posit additional fundamental properties, since an understanding of the way things are physically gives us all we need to understand how things are phenomenally. So, whilst one form of dualism might yield a theory of consciousness, we would be unable to grasp why this theory should be true. By my lights, such a theory thus falls short of an *explanation* of consciousness⁹.

A similar problem arises for type-B materialism. To see this, consider Jackson's (2003) discussion of the reduction of facts about temperature in gases to facts about mean kinetic energy in gases. This reductive explanation will be of the following form:

- 1) Temperature in gases is that which does so and so (*a priori* premise about the concept of temperature)
 - 2) That which does so and so is mean molecular kinetic energy (empirical premise)
 - C) Therefore, temperature in gases is mean molecular kinetic energy.
- (Jackson (2003), p.5)

Jackson argues that if we are to identify temperature in gases with mean molecular kinetic energy, we need the first *a priori* premise in order to avoid the possibility of our being eliminativists about temperature in gases. For without that connection between temperature and the functional role ('that which does so and so') which premise 2) tells us is occupied by mean molecular kinetic energy, there is nothing to prevent our explanation being interpreted as explaining *away* temperature – we thought there was such a property as temperature, though actually there is only mean molecular kinetic energy. The *a priori* premise rules this out, since it tells us how our concept of temperature is related to the functional role occupied by mean molecular kinetic energy. So, without the right sort of conceptual mesh between the domains in our explanation, an explanation of consciousness risks being indistinguishable from an *elimination* of consciousness, asserting that the closest thing there is to consciousness is a functional role we have identified, but admitting that

⁹ See Dennett's (1993) reply to Foster for similar arguments for the explanatory impotence of dualism.

there is no conceptual link between this functional role and consciousness conceived in the first-person, phenomenal way we were interested in explaining.

Chalmers (2003) criticises explanations appealing to *a posteriori* identities or necessitation in a more general way. If the functional description and phenomenal target are held to be identical by our explanation, then it seems that this identity will be “epistemically primitive” (p.15), since it will not be an identity that we could deduce just by considering the functional and phenomenal facts – if such a deduction were possible, then we would have an explanatory story about how and why our functional description is one and the same as our phenomenal target, giving us the *a priori* reduction we were after, and making our position a type-A materialism. Likewise, if the functional is said to necessitate the phenomenal without being identical to it, this will be what Chalmers calls a “strong necessity” (p.16) – a necessity we posit despite being unable to see how or why one level necessitates the other.

Chalmers wants to conclude that these considerations militate in favour of dualism. Because we can't give an explanation of the relevant identity or necessitation in terms of the mesh between the concepts in the two domains, the identity or necessitation we posit should be viewed as a fundamental law in our explanatory theory. But if our explanation of consciousness is in terms of fundamental laws linking domains that are at least conceptually distinct, then it looks like we end up with a kind of dualism. Even if someone holding such a position wanted to remain a physicalist, they would be a physicalist who claimed that there are (at least) two distinct types of physical property (the functionally specifiable ones and the phenomenal ones) and that statements about one of these physical domains cannot be translated into statements about the other – rather the two domains are linked by contingent, empirically discoverable laws. Such a position seems to be physicalism in name only.

Coupled with Jackson's argument, above, we might think that this shows that *a posteriori* explanations of consciousness look to be neutral between eliminativism, physicalism and dualism. Accepting these arguments implies that the physicalist must demonstrate a conceptual mesh between the phenomenal and the physical if their theory of consciousness is to be logically different from the non-explanatory positions of eliminativism and dualism. As Chalmers (1997, p.10) puts it, “For a satisfactory theory, we need to know more than *which* processes give rise to experience; we need an account of why and how. A full theory of consciousness must build an explanatory bridge.”

Endorsing the conceptual mesh requirement apparently commits us to type-A materialism – since it holds that there is an *a priori* entailment between the phenomenal and the physical, it is the only position above which holds that the relation between the two

domains is one that we can grasp in the necessary way. However, it seems difficult to reconcile the intuitions pumped by the arguments of the previous section with the possibility of type-A materialism's truth. If reflection on our experience allowed us to see that its physical properties were all that required explaining, type-A materialism would be straightforwardly vindicated – we would know that a conceptual mesh obtained between the phenomenal and the physical since experience told us that phenomenal properties are instances of physical properties. But the intuition behind the knowledge argument and the apparent conceivability of zombies and inverted spectra show that this is not what reflection on our experience tells us. Rather, it seems to be an experiential datum that experience has phenomenal properties that go beyond the structural, functional and intentional properties to which physicalism can appeal. By asserting the presence of a conceptual mesh between physical properties and the felt properties of our experiences, type-A materialism apparently commits itself to denying this datum, and thus to an unpalatable eliminativism about the phenomenal properties we wished to explain. The type-A materialist might try to provide arguments that introspection reveals only physical properties, but it is hard to see how such arguments could overcome the force of our intuitions to the contrary¹⁰.

I have argued that none of the positions considered in this section possess the resources for a satisfactory explanation of consciousness. There are various ways we might try to resist this conclusion. In defence of type-A materialism, we might argue that our conceptual framework should be answerable to our best empirical theory, and thus that our concepts of phenomenal consciousness should be physical ones picked out by neuroscience, allowing an *a priori* reduction of the phenomenal to the physical (Lamme (2006)). In defence of type-B materialism, we might argue that our claims of identities or necessitations between the physical and the phenomenal should not be seen as fundamental posits, but as instances of inference to the best explanation warranted by the perfect correlation we will hopefully discover between the physical and the phenomenal (e.g. Churchland (2005), Hill (1997)). More generally, we might question the conceptual mesh requirement by casting doubt on whether paradigm instances of reductive explanation really have an *a priori* component (Block and Stalnaker (1999)). More generally still, we might give reasons for general skepticism about the *a priori*, and conclude that the entailments being demanded do not exist (e.g. Quine (1951), Laurence and Margolis (2003)).

Whilst these are all interesting avenues, I don't think we need pursue any of them until more orthodox responses have been properly considered. I propose to concede to

¹⁰ See Chalmers (2003) Section 4 for a survey and critique of current candidates for such arguments.

proponents of the explanatory gap both that explanations of consciousness drawing on *a posteriori* necessities leave explanatory questions that we might care about unanswered, and that a conceptual mesh obtaining between the domains in our explanation is a desirable feature of a reductive explanation. I think it's worthwhile seeing how far we can get whilst accepting these constraints. If the account we arrive at does not meet these constraints, we should be able to say something about why they should not apply in the specific case of phenomenal consciousness. If it turns out that we cannot meet these constraints, or say why they should not apply in this case, we can return to the more general lines of resistance above. Part of the task of this thesis is to argue that this will not be necessary. The next section presents the explanatory model occupied by the theory I will outline and defend.

1.3 Deflationism Defined and Defended

The claim that an explanatory gap exists has been motivated by asking us to reflect on our experience, and claiming that what we find will not seem amenable to the sort of reductive explanation that proponents of the explanatory gap demand of physicalists. The last section argued both that we should (at least temporarily) grant these points about experience and explanation, and that none of the theories in Chalmers' taxonomy afford good prospects for a suitable explanation. However, in this section I argue that between the unsatisfactory type-A and type-B materialisms, there is space for a promising solution to the problem of the explanatory gap.

It is the reliance on intuitive conclusions about the deliverances of introspection that gives us the necessary room to manoeuvre. It is open to the reductionist to argue that although we're prone to think we're acquainted in experience with properties that are not physically analyzable, this is a mistaken characterisation of our introspective evidence. There certainly seem to be cases where subjects are prone to misdescribing their experience – for example, we have a tendency to assume that our colour vision is uniformly sensitive from the centre of the visual field to its periphery, but proper scrutiny of the way things seem to us, perhaps with the help of visual aids such as a coloured playing card held at the periphery, will reveal that this is not the case (Dennett (1991))¹¹. Perhaps we make a similar mistake when we think we find physically unanalysable properties of our experiences. Alternatively, we might grant that introspection really tells us that there

¹¹ See also Schwitzgebel (2002, 2008) for a variety of cases where subjects appear to have poor knowledge of the character of their conscious experience.

are such properties, but try to explain how introspection misleads us here – we might think that if we give an explanation of how and why introspection distorts how things are with experience, and show that a non-distorted characterisation of experience can be reductively explained, then this constitutes an explanation of all the puzzling features of experience¹².

We can see each of these options as **deflationist**, since they both aim to deflate the apparent introspective evidence which motivates the positing of an explanatory gap, rather than taking that evidence at face value¹³. If we grant the above requirements on explanation, and that thinking about our experience makes it tempting to posit an explanatory gap, then our reductive explanation must be deflationist in this way, since alternative construals of our introspective evidence involve retaining a physically unanalysable explanandum. Such a position does not fit neatly into the mould of either type-A or B materialism as defined by Chalmers, and avoids the features of each that made them unsatisfactory. Unlike type-B materialism, our deflationist alternative does not require explanatorily primitive principles linking the physical and the phenomenal – the deflationist claims that once we construe the introspective evidence correctly¹⁴, we will be able to grasp how the phenomenal properties of our experience reduce smoothly to the physical. Unlike type-A materialism, this reduction does not come at the price of denying that naïve introspection presents us with evidence for phenomenal properties of experience that are irreducible to physical properties – the deflationist claims either that we can distinguish between informed and naïve introspection, and that only the latter

¹² Setting up a dichotomy in this way looks like it presumes a strong realism about the deliverances of introspection, and a corresponding sharp divide between our experience and our acquaintance with it. We might not be happy with either of these. The position I defend is neutral between these two deflationist options, and I take this as a point in its favour, since I'm suspicious of the conception of experience required to see the distinction between them as a substantive one. My reasons for this emerge as I develop my account, particularly in chapter 6.

¹³ Such a construal of deflationism might do it a disservice, since we should question whether the interpretation of the introspective evidence required to motivate the explanatory gap is really the natural one. Terminology with the opposite rhetorical bias might call proponents of the explanatory gap 'inflationists', and deflationists 'naturalists'.

¹⁴ For the first type of deflationism mentioned in the previous paragraph, this will involve coming to see that our introspective evidence is other than we initially took it to be. For the second, this will involve coming to see that our introspective evidence misleads us as to the sorts of properties our experience has. Recall from footnote 12 that I am skeptical about the difference between these two options, and my eventual account (see chapter 6) does not distinguish between them.

seems to provide evidence for irreducible phenomenal properties, or that introspection does present us with such evidence, but misleads us in doing so¹⁵.

An analogy will help clarify the position being outlined. We see a magician place a ball under a cup, then tap the cup with a wand. When she lifts the cup, two balls are underneath. She replaces the cup, taps it then lifts it again, and the balls have disappeared. How should this trick be explained? Taking our description at face value seems to commit us to events in the world that are not readily explicable in physicalist terms – objects appearing from nowhere then vanishing into thin air, or perhaps being teleported. A type-C materialist about the trick will argue that vanishings and teleportations are explicable in physical terms, but that we are either currently or necessarily unable to see how such an explanation will go. A dualist about the trick will argue that we should expand our physicalist ontology to make room for fundamental sorts of events such as vanishings and teleportations. A type-B materialist will argue that vanishings and teleportations are identical to physical events, but that there is no conceptual mesh between their properties as vanishings and teleportations and their properties as physical events that lets us deduce one set of properties *a priori* from the other. A type-A materialist will argue that nothing in what we have seen resists explanation in physicalist terms, either by claiming that there is an *a priori* path from physical properties of events to their properties as vanishings and teleportations, or that the trick contained no vanishings or teleportations to be explained in the first place.

None of the outlined explanations are adequate, for reasons we saw in the previous section. The first constitutes a denial of the possibility of an explanation, the second stipulates that the sorts of events for which we wanted an explanation are fundamental features of reality, and the third stipulates that there is an identity between explananda and explanans, but claims that we cannot grasp why this identity holds. The fourth owes us either a story about how we can move *a priori* from the problematic explananda to the physical explanans, or why it is legitimate to deny that apparently manifest explananda exist. The deflationist position lies between our third and fourth options. The deflationist argues that thinking more about how the trick could be turned lets us see that positing teleportations and vanishings is not the only way we could explain what we have seen.

¹⁵ Focussing on the first option might prompt us to categorise deflationism as a type-A materialism, since the phenomenal properties of our experiences *suitably construed* are *a priori* reducible to the physical. Chalmers does not explicitly address whether type-A materialists are required to provide an *a priori* reduction of the naïve construal of phenomenal properties, and I think answering the question of whether such a deflationism qualifies as a type-A materialism ends up being a stipulative matter. Note, though, that in chapter 6 I argue that my deflationist position is usefully classified as a sort of phenomenal concept strategy, a position Chalmers classes as type-B materialist, supporting my contention that my position falls between the gaps of Chalmers taxonomy.

Empirical and conceptual investigation reveals that the description of the trick on the previous page does not commit us to the existence of teleportations and vanishings construed as such – rather, the facts as we have described them are explicable via an appeal to sleight-of-hand, a phenomenon which is *a priori* reducible to the physical. The deflationist explanation acknowledges the type-B materialist’s conviction that our evidence makes it natural to posit the existence of teleportations and vanishings, and agrees with them that a satisfactory explanation of the trick should make reference to such events. They also agree that, understood under the description of teleportations and vanishings, such events are not reductively explicable. But, they claim, investigation reveals that such a description is not the best way to make sense of our evidence. Once we understand how the trick is done, we can identify vanishings and teleportations with reductively explicable sequences of sleight-of hand manoeuvres. The deflationist explanation of the trick, then, includes reference to vanishings and teleportations and an acknowledgement of our reasons for believing in them, but shows us how and why we should redescribe those events in a way that renders them reductively explicable.

The crucial feature of the trick for the deflationist is that the evidence we are attempting to explain forces us to make reference to vanishings and teleportations, but does not force a conception of those events upon us that renders them reductively inexplicable. The deflationist about phenomenal properties argues that the same is true of our introspective evidence. Reflecting on our experience gives us reason to posit reductively inexplicable phenomenal properties of experience, but empirical and conceptual investigation can reveal how and why we should redescribe these properties in a way that renders them *a priori* reducible to the physical. The crucial feature of our introspective evidence for the deflationist is that it dictates that there are phenomenal properties of experience, but not how those phenomenal properties are described. If the deflationist can give an account of how those properties can be plausibly redescribed in such a way as to render them reductively explicable, they will have provided a solution to the explanatory gap that avoids the pitfalls of type-A and type-B materialism¹⁶.

I have argued that this deflationary type of explanation occupies a place in between type-A and type-B materialism, and avoids the problematic features of each. Might Chalmers, or someone with a similar philosophical disposition, object to such a strategy? It seems possible that Chalmers might think the position fails to “take consciousness seriously”¹⁷,

¹⁶ A similar conclusion is argued for by Carruthers (2004). Lewis (1995) is also suggestive of such a position.

¹⁷ See the introduction to Chalmers (1996)

by ultimately understanding experience as a physical phenomenon, in spite of apparent introspective evidence to the contrary. Whilst I agreed above that this might be an appropriate objection to type-A materialism, I think that it begs the question in the present context. The deflationist, like Chalmers, starts with introspectively known phenomenal properties as her explanandum, and, like Chalmers, acknowledges the temptation to construe these properties in a way that makes them inaccessible to physicalist explanation. Thus, she cannot be justifiably accused of trivialising the problem of consciousness by ignoring or denying the explanandum. Where she differs from Chalmers is in denying that the anti-physicalist construal of these properties is the only one available, though she does not claim that the alternative construal is itself transparently available to introspection¹⁸. If Chalmers is to object to this sort of strategy, he must claim that the introspective evidence admits only one metaphysical gloss on the phenomenal properties of experience – and this doesn't seem to be something we can settle by introspection alone. It may be that all other glosses offered are somehow untenable, and we are thus forced to an anti-physicalist conception of experience. But in this case, the candidate glosses must be assessed on their merits. It is the task of this thesis to provide such a gloss, and offer empirical and conceptual arguments in its favour.

1.4 The Duties of the Deflationist

This lets us see what is required of a deflationary explanation of phenomenal consciousness. It must accommodate the introspective evidence that leads to the putative explanatory gap, but plausibly redescribe that evidence in a way that renders it compatible with physicalism. The redescription must be such as to explain why the anti-physicalist construal of the evidence is natural, but mistaken, and why the deflationist construal is not introspectively obvious. Additionally, the deflationist construal of experience must be empirically plausible – otherwise the deflationist has shown only how the explanatory gap *could* be closed (perhaps a worthwhile achievement in itself), but not how it can be closed for *our* experience.

Thus far in setting out the problem posed by phenomenal consciousness for reductive explanation we have supposed that introspection makes it obvious that phenomenal properties pose a problem for physicalism. Whilst I think this simple appeal to introspection is sufficient to make the case for an explanatory gap between the phenomenal

¹⁸ Instead, it must be developed by considering the kinds of arguments and data I set out in subsequent chapters of this thesis, particularly chapter six.

and the physical, some go further, and attempt to spell out the features of phenomenal properties that make them reductively problematic.

Firstly, phenomenal properties seem to be **ineffable** - the way things seem for the subject cannot be precisely expressed in words. We might be able to allude or gesture to the phenomenal properties of our experience, but any amount of metaphor or verbal description will necessarily fail to characterise them completely. No matter how eloquently I describe the taste of the beer I am drinking, or the colour of the sky I am viewing, my description will always lack some of the richness and particularity of my experience. Relatedly, phenomenal properties seem to be **private** - they belong and can be fully known only to the subject who instantiates them. Another feature of our acquaintance with the phenomenal properties of our experience is that they are **directly available** to us in a way that (for example) the neural properties of experience are not. Some make the apparently stronger claim that our knowledge of phenomenal properties is incorrigible. We might wish to say that there are cases where we can be mistaken about the character of our own experience, as suggested in 1.3¹⁹, and so deny that such knowledge is incorrigible in a strong sense. However, it seems plausible that there is a weaker sense in which such knowledge is incorrigible - although we can make mistakes in the way we construe or describe our experience, we cannot be mistaken about the way the experience seems to us. Suppose I sincerely but mistakenly claim to be sober on the basis of introspection. The next day, if I can bring my experiences back to mind with sufficient acuity, I might realise that this was not in fact what introspection was telling me - it really did seem as if I was drunk, but I made an erroneous judgement or report on the basis of that seeming. Finally, it is sometimes claimed that phenomenal properties are **intrinsic** properties of experience²⁰. Intrinsic properties are notoriously difficult to define precisely, but presumably the idea here is that the way phenomenal properties of an experience are not dependent for their existence on factors other than that experience itself, such as the dispositions or abilities that experience gives rise to²¹.

A further requirement on an adequate deflationist theory, then, is that it explain the compatibility of these putative properties of experience with physicalism. However, the

¹⁹ See Schwitzgebel, *op cit*. Someone who mistakenly believes that they are sober on the basis of introspective evidence might be a good example of such a case.

²⁰ See (e.g.) Harman (1990), and references. Harman argues against the thesis that phenomenal properties are intrinsic properties of experience.

²¹ This list of problematic properties is adapted from Dennett (1988). I have omitted other properties which a proponent of irreducible qualia might wish to include on the list where I think they are entailed by properties I have included, or have been unable to find a sufficiently clear formulation of the property being claimed.

deflationist need not concede that these are all genuine experiential properties. Since they aim to explain introspective evidence, but attempt to motivate a particular construal of that evidence, it is open to them to argue that we attribute one or more of these properties to our experience in error. It is sufficient for the deflationist to explain merely why it is natural to construe our experience as possessing these properties, without granting that the properties actually obtain. This is particularly important for the above property of intrinsicity - claiming that the phenomenal properties of an experience are intrinsic to it entails that those properties do not depend on anything external to that experience for their existence. Whilst this may be compatible with those properties being reductively explicable, it places a strict limit on the domain in terms of which we might explain them. Deflationists can eschew this limit if they can provide an argument that the introspective evidence makes natural but does not *force* a construal of phenomenal properties as intrinsic upon us. I won't be in a position to make a full case for how I think a deflationist should treat these properties until chapter six, so I defer further discussion of them until then. In the meantime, let's start to equip ourselves with some of the tools we will need to construct our deflationist position.

1.5 Two Enactive Theories

There is a growing consensus that an adequate philosophical understanding of conscious perception must make essential reference to its links to skilled worldly action²². We can call a view that makes such an essential appeal to action an *enactive* view. The 'essentiality' of the appeal indicates that the importance of action is in virtue of more than its role in effecting some other (perhaps neuronal or cognitive) properties in terms of which consciousness is to be explained - action is being given a *non-instrumental* role in the account of perception. One motivation for such an enactive focus might be the intuition that whenever we find an agent with the capacity for flexible and intelligent action, we intuitively deem it conscious²³. Another might be the adaptationist conviction that consciousness would not exist unless it were of some value to its possessors, and that this value must consist in the ability of conscious entities to act in a way that increases their chance of surviving and reproducing - thus comparing the sort of things that conscious creatures can do with those that intuitively non-conscious creatures can do might provide

²² See e.g. Varela, Thompson and Rosch (1991), Hurley (1998), Noë and O'Regan (2001), Noë (2004), Matthen (2005), Schellenberg (2007), Thompson (2007).

²³ This intuition seems less strong in the case of 'artificial' agents, such as sophisticated robots, prompting Hurley (1998, ch.4) to suggest that *life* is required for the intuition that consciousness is present to take hold.

an account of the function of consciousness, and perhaps its identification with this functional role. I won't try and fully motivate adopting such a focus in this section, however. This thesis will argue for a particular brand of enactivism, citing empirical evidence in its favour and arguing that it can provide a way of fleshing out the kind of deflationary explanation of conscious experience argued for above. A key part of my strategy will be contrasting the 'action-space' account which I favour with the 'sensorimotor' account of Noë (2004), in order to argue against Noë's position and to clarify and motivate my own. The character and motivations of the views will emerge in more detail over the following chapters of the thesis - my aim in this section is merely to set out two competing views of the relationship between action and perception, rather than fully motivate either position.

According to the sensorimotor theory of Noë (2002, 2004, 2005), the content and character of our visual experience is determined by our implicit knowledge of the systematic ways in which sensation will change as a result of certain bodily movements²⁴. Perception, on this view, is not merely a matter of passive reception of sensory stimulation. Rather, the perceiver is partially in control of the sensory information she receives, since this varies with her movements. Perception acquires its content and qualitative character through the perceiver's acquired knowledge of this interdependence; through the perceiver understanding movement-dependent patterns, regularities and contingencies in the sensory stimulation she receives. For example, a line in front of the perceiver appears vertical just in case she deploys her implicit knowledge that her sensations will remain largely the same if she nods her head up and down the line, but will differ in a predictable and regular way if she moves her head from side to side. A visually-presented tomato appears spherical (rather than appearing as a circular tomato-façade) if the perceiver possesses implicit knowledge of how her sensations would change were she to move around it. And the tomato is experienced as visually rather than tactually presented if the perceiver implicitly knows that (for example) moving her head and eyes around will alter her sensations in characteristic ways, but leave her tactile sensations unchanged.

Much of the support for Noë's position comes from discussion of various empirical work. To take a single example, consider Bach-y-Rita's work with tactile-vision substitution

²⁴ Sensorimotor views are also defended by Hurley (1998), Schellenberg (2007) and O'Regan and Myin (2002). I briefly discuss some of the other forms sensorimotor theory may take in the next chapter, and discuss Hurley's position in more detail in chapters 5 and 6. I focus chiefly on Noë's view in this thesis, since his position provides the best contrast with the action-space account I develop.

systems (TVSS)²⁵. Such devices consist of a head-mounted camera which transduces patterns of optical information into isomorphic patterns in an array of vibrating rods on the back, abdomen or thigh of a blind subject. After a short period wearing the device, subjects no longer perceive the signals from the device as vibrations on their body – rather, they perceive them as originating from distal objects, out there in the world. They thus begin to act and speak as if they are in visual contact with the world, ducking when objects loom at them, describing occluded objects as hidden from view, and making the kind of correct judgements about the size, shape and position of objects in her environment that we would expect of a sighted subject. According to the sensorimotor view, this is because they have begun to visually experience their environment, and this in turn is because the TVSS device provides an approximation of the relations which obtain between visual perceptual stimulation and movement for the subject. Importantly, TVSS-users must go through a period of actively exploring their environment whilst wearing the device, thus implicitly learning about the relations between stimulation from the device, their own movement, and objects in their environment. Thus, claims the sensorimotor view, to the extent to which TVSS replicates these relations, users will enjoy visual experience.

All perception, according to the sensorimotor view, is a matter of implicit understanding of relations between sensation and movement (sensorimotor contingencies) such as these, and implicit expectancies concerning the ways in which sensation would vary with the perceiver's movement, according to these contingencies. Thus, perceiving on Noë's account is a matter of *knowing how what we can do affects what we can see*.

The dispositional account of Pettit (2003) also emphasises the role of skills and implicit understanding in perception, but in a different way to the sensorimotor view. For Pettit, for something to look a certain way is for it to empower certain abilities in the perceiver. On this view, a tomato's looking red to a perceiver is a matter of its empowering her to, among other things, sort it with red and other similarly-coloured objects, sift it from differently-coloured objects, or different shades of red, and track it across a range of different backgrounds and perceptual situations. Once we have detailed the various sifting, sorting and tracking abilities enabled in a perceiver by exposure to a coloured object, we have said all there is to say about her experience of that colour²⁶.

²⁵ See Bach-y-Rita and Kerzel (2003) for a recent review.

²⁶ I develop and defend a related view of colour perception in detail in chapters 3 and 4, so will not try to motivate this view further here.

Though Pettit restricts his treatment to colour looks, his account might be generalised to other aspects of perception. For example, Matthen (2005) extends this kind of approach to sensory perception in general, arguing that:

“[P]erception is for action of two kinds. First, it is for the guidance of the body as it interacts with other material objects. Second, it is for finding out about things in the world, for building up a record of the characteristics of such objects, and forming expectations concerning how they should behave in the future. The main thesis to be advanced builds the content of perception on what it tells us about how objects should be treated with respect to these goals.” (p.8)

It might be claimed, for example, that a tomato’s looking spherical to the perceiver is a matter of her perception of it disposing her to sort it with other spherical objects and sift it from differently shaped ones. And the tomato is experienced as visually, rather than tactually presented just in case it empowers a suite of abilities in the perceiver that are characteristic of vision rather than touch (sifting and sorting it on the basis of its colour, rather than, say, its temperature to the touch).

An important feature of the view to be defended is that the enabled abilities in terms of which perception is to be explained must be available for integration with the subject’s capacities for practical reasoning and action-planning (Clark (2006b, 2007)). My account of just what this requirement amounts to will be given in chapter 5, and we will see empirical evidence in its favour in the next chapter. For now, I’ll merely note that the basic motivation for such a requirement is that perceptually enabled abilities seem entirely consistent with the absence of conscious experience – a simple robot can use sensory information to guide its behaviour without, presumably, being a seat of perceptual consciousness. As we shall see in chapter 5, the action-space theorist argues that this is due to the unavailability of that sensory information to systems of practical reasoning and action-planning – rather, that information is automatically and inflexibly used to guide behaviour, thus failing to enable a conscious experience according to the action-space account’s criteria.

As with the sensorimotor theory, proponents of this view of the relationship between perception and action appeal to various strands of empirical work to support their position. Pettit appeals to results showing adaptation to colour-distorting goggles (which we will discuss in detail in chapters 3 and 4), and Matthen appeals, amongst other things, to the conception of sensory processing at work in our most successful current psychological

theories. The remainder of this thesis tries to motivate such a view, so I shall not attempt to do so here. For present purposes, note only that Pettit and Matthen both attempt to explain perception in terms of the ways that a perceiver's sensitivity to her environment enables her to act. We can call a theory that appeals to action in this way an 'action-space' account. Perceiving, according to such an account, is a matter of *knowing how what we can see affects what we can do*.

1.6 Enactivism and Deflationism

So, we have two contrasting ways of understanding the relationship between action and perception. The camp occupied by Noë thinks we must appeal to action in understanding perception since perception is constituted by our understanding of how possible perceptions depend on what we might do. The camp occupied by Pettit, Matthen, and myself, argue for an opposite direction of dependence between perception and action, claiming that perceptual experience is constituted by our understanding of how possible actions depend on what we perceptually detect. What makes each of these positions a potential deflationary solution to the explanatory gap?

Sensorimotor theorists often stop short of addressing the explanatory gap itself, claiming that an appeal to sensorimotor contingencies can provide an account of the properties in terms of which one experience differs from another, but remaining silent on whether such an appeal might explain why a subject enjoys some experience rather than none²⁷. However, in their bolder moments, they sometimes claim that their enactive focus reconfigures the explanatory problem, and lets us see how it may be addressed. For example, Noë and O'Regan (2001) state that:

“In our view, the qualia debate rests on what Ryle called a category mistake. Qualia are meant to be properties of experiential states or events. But experiences, we have argued, are not states. They are ways of acting. They are things that we do. [...] Qualia are an illusion, and the explanatory gap is no real gap at all.”
(p.960)

Noë and O'Regan take care to emphasise that they are not denying that there is something-it's-like to perceive various aspects of the world. Rather, they deny that we

²⁷ See (e.g.) Hurley and Noë (2003), Noë (2004), p.230-1. Hurley and Noë thus distinguish between a 'comparative' explanatory gap, which they attempt to address, and an 'absolute' explanatory gap, about which they remain largely silent. I criticise this position in chapter 5.6.

should understand this manifest feature of experience in terms of our acquaintance with inner states or events, such as 'qualia' traditionally construed. Rather, experiences are temporally extended encounters with our environment, more akin to processes than states or events. We could fit this position into our deflationist schema by interpreting Noë and O'Regan as claiming that introspection makes it tempting to posit physically inexplicable experiential states where there are none, but when we view our introspective evidence accurately, we see that it is compatible with experience being a matter of active engagement with the world, rather than acquaintance with a qualitative state. A sensorimotor deflationist would then argue that such a conception of experience makes it amenable to physicalist explanation, and cite the empirical work implicating grasp of sensorimotor contingencies in perception as evidence that their theory was true of human conscious perceivers. Lastly, they must show that their conception of experiences as grasp of sensorimotor contingencies has the resources to explain why phenomenal properties are experienced as ineffable, private, directly-available and intrinsic.

Whilst Matthen also leaves questions about the explanatory gap largely to one side²⁸, such questions are a major part of the focus of Pettit's account. Pettit offers his account as an antidote to the anti-physicalist account of colour experience as involving acquaintance with an intrinsic, effect-independent property of experience. We should resist this characterisation of our introspective evidence, in favour of one according to which our experience of an object's colour is wholly explained by the abilities enabled by our perceptual exposure to it. Thus,

"The way the object looks has no character that we might imagine being dissociated from the enabling effects, being present in their absence, or absent in their presence. It has a character only as something that manifestly enables the subject to sift, sort and track in a certain way, and thus to make corresponding judgements; it has a character that is essentially tied to those empowering effects." (Ibid. p.229)

This construal of our experience is not introspectively obvious, since "...we come to conceive of red, and we come to see red *through* its dispositional aspect, if not *under* it." (p.225) – we don't need to see redness *as* a property that disposes us to sift, sort and track objects in various ways for that to be the correct account of that property. We might explain this fact by arguing that these dispositions come about as the result of subpersonal processing of whatever features subserve the objects colour. So,

²⁸ Except briefly in its comparative form ((2005), p.241-2).

“The idea is that perceptual exposure to a red object tunes us in a way that enables us to sift, sort and track in various ways, and that this tuning occurs ‘behind our back’; it occurs in virtue of the brain’s response to the perceptual input at the subpersonal level. What we become aware of in the experience presupposes that those subpersonally induced responses are in place; the red look of the object is the look it has so far as the perceived object is registered as enabling – as having already enabled – those responses: so far as it is registered as having that active power.” (p.232)²⁹

The idea is that such a story could be fleshed out to provide both a physicalist-friendly construal of our introspective evidence, and an explanation of why such a construal is not immediately evident to us. The above quotes from Pettit already suggest his rejection of intrinsicity as a property of experience – he argues that rather than our colour-categorising abilities being enabled by effect-independent properties of experience, those enabled abilities help constitute that experience. The hope is that charting the enabled abilities essential to colour perception will also yield explanations of the other putatively problematic properties of experience detailed in section 3, and that this conception of visual perception is empirically justifiable.

So, these are the bare bones of deflationist sensorimotor and action-space theories of phenomenal consciousness. The remainder of this thesis will be devoted to arguing that the action-space theory should be preferred over its sensorimotor counterpart, both for its better mesh with empirical work and its ability to meet the constraints on an adequate deflationary theory of experience that we have set out in this chapter.

Conclusion

In this chapter, we have looked at the motivations for believing that certain properties of our perceptual experience pose a problem for physicalism. I argued that we should see the challenge posed by the explanatory gap as a demand for an explanation of phenomenal consciousness which provides a conceptual mesh between experience and a physicalist-friendly explanans. Once we see the challenge in this way, we see that none of the positions surveyed in section 1.2 look tenable. However, I argued that there is room for an intermediate, deflationary position, between type-A and type-B materialism. According to such a position, we can redescribe our introspective evidence in a way that renders it

²⁹ Pettit intends such ‘registration’ to be a subpersonal-level phenomenon.

compatible with physicalism, whilst explaining the temptation to construe it in an anti-physicalist way. To be plausible, such a position must do justice to our introspective evidence, and to the ease with which it can be used to generate anti-physical conclusions. It must also be empirically plausible, and provide explanations of putative problematic features of experience such as intrinsicity, ineffability, privacy and direct-availability. We briefly surveyed two enactive theories of perception with a structure that might permit them to discharge these obligations – Noë's sensorimotor theory, according to which conscious perception is a matter of understanding the likely sensory effects of movement, and the 'action-space' theory, according to which conscious perception is a matter of understanding the ways one is poised to act on the basis of one's perceptual sensitivity to the environment.

Chapter 2: Action, Perception and the Two Streams

In chapter one, we set out the problem of the ‘explanatory gap’ between the phenomenal and the physical and the constraints on a satisfactory solution, then raised the possibility that an enactive theory of consciousness might have the structure necessary to meet these constraints. The two enactive theories we briefly considered were Noë’s sensorimotor theory, according to which visual experience arises through a grasp of the perceptual consequences of movement, and the ‘action-space’ theory, according to which visual experience arises through a grasp of the practical consequences of perception. The next three chapters are chiefly devoted to adjudicating between these views, and ultimately finding in favour of the relationship between action and perception advocated by the action space theorist. In chapters three and four we will see that the action-space account has resources to give an account of colour experience that are unavailable to the sensorimotor theorist. In this chapter, I want to consider what the dual visual systems hypothesis (DVS) - a currently influential theory in the neuropsychology of vision - suggests about the functional underpinnings of conscious vision, and whether it militates in favour of one or the other of our enactive views. I argue that, contrary to the views of several recent commentators, this hypothesis is not directly inconsistent with a sensorimotor theory of visual perception. However, whilst DVS is not strictly incompatible with the sensorimotor theory, I argue that it suggests a view of the nature of conscious visual experience that has strong affinities with the action space approach, and according to which knowledge of sensorimotor contingencies plays, at best, an indirect role in visual perception. We should conclude that our current best neuropsychology of vision gives us reason to favour the action-space approach over the sensorimotor theory, and that we thus have empirical reason to focus on the action-space approach in our attempt to develop a deflationary theory of phenomenal consciousness.

In section one, I review the commitments of various possible sensorimotor theories of perception. Section two is an overview of the features of the dual visual systems hypothesis relevant to our consideration of the debate between action-space and sensorimotor theories. Section three considers various ways in which DVS has been used to challenge Noë’s sensorimotor view, and argues that Noë and sensorimotor theory can respond to each. However, in section four I argue that the real challenge from DVS to sensorimotor theory is the affinity that obtains between the action-space account and DVS’s suggestions about the functions of conscious vision. This suggests that the importance of sensorimotor contingencies should be relegated to (at best) a possible implementation of this functional role, thus compromising the sensorimotor theorist’s claim that an appeal to

sensorimotor dynamics is of non-instrumental importance in an account of perceptual experience. Lastly, in section five I address the challenge that the standard conclusions from DVS to the functional role of conscious vision are too hastily drawn. I conclude that DVS gives us empirical reason to favour an action-space account over a sensorimotor view in developing a deflationary theory of phenomenal consciousness.

2.1 Sensorimotor Theory Reviewed

Recall from the previous chapter that a sensorimotor theory of conscious perception is one that holds that the content and character of perception is determined by a perceiver's understanding of the ways in which possible actions will effect possible perceptions. Crucially, sensorimotor theorists argue that such sensorimotor dependencies are of *non-instrumental* importance to the generation and explanation of perception. Everyone can agree that possible movements are important to the content and character of a perceiver's experience, since her movements will usually alter what she perceives. But the sensorimotor theorist's view of the relationship between action and perception goes beyond a bland appeal to actions as means to different perceptions. On their view, a perceiver's grasp of possible movements and the perceptions that they would result in is constitutive of the content and character of experience. The discussion of TVSS in 1.5 illustrated this claim – according to the sensorimotor theorist, the 'visual' experiences of TVSS users are like the visual experiences of normally sighted perceivers to the extent that both involve a grasp of an isomorphic set of contingencies between possible movements and possible patterns of perceptual stimulation.

This sketch of sensorimotor theory leaves room for several possible variants, due to the ambiguities in the notion of dependence between perception and action. The senses of 'perception' and 'action' being appealed to can be cashed out in a number of ways, and different views of the sensorimotor contingencies to be appealed to are thus possible. For example, 'perception' admits of both a personal-level construal, denoting contentful experience, and a subpersonal construal, denoting patterns of retinal stimulation, or subsequent neural activity³⁰. 'Action' in this context can likewise be given a personal-level

³⁰ This ambiguity can be used to create a potential dilemma for the sensorimotor theory taken as an attempt to explain phenomenal consciousness. For if perceptual stimulation in the contingencies appealed to is construed subpersonally, we are owed a supplementary story about why expectations about subpersonal stimulation should give rise to experience. And if the appeal is to contingencies between movement and personal-level perceptual experience, then the theory presupposes the phenomenology it is intended to explain. This dilemma has been independently noted by several critics of the sensorimotor view, including Block (2005) and Clark (2006), and I won't pursue it further here.

construal, denoting gross intentional actions aimed at the fulfilment of a goal of the agent, or a subpersonal construal, perhaps denoting efferent neural signals corresponding to motor commands. Intermediate construals of 'action' also seem possible – exploratory eye movements are bodily actions of the agent, but there is good reason to believe that they are not always under intentional control³¹. Each construal of perception can be combined with each construal of action to yield a different sort of sensorimotor contingency, and sensorimotor theorists might appeal to some or all of these in their account of perception. For example, Noë and O'Regan (2001) focus chiefly on the effects of exploratory head and eye movements on subpersonal retinal stimulation, citing (for example) the way in which the retinal stimulation from a vertical line will remain largely unchanged as a perceiver nods her head up and down, but change in a regular and predictable way as she moves her head from side to side. Noë (2004) focuses chiefly on the relationship between such exploratory movements and perceptual experience, discussing (for example) the changing ways a plate or a penny *looks* to the perceiver as she moves in relation to it³². An appeal to contingencies between higher levels of perception and action is at work in O'Regan Myin and Noë's (2002) suggestions that the feeling of material wealth might be partially explained by the knowledge of the likely consequences of the intentional action of attempting to withdraw a large sum of money from a cash machine, or that the feeling of confidence might be a matter of expecting the successful consequence of some intentional endeavour. And Hurley (1998, ch.9) appeals to the relations between the lowest levels of perception and 'action' in her explanation of Gallistel's (1990) report that the visual scene seems to move as a result of the attempted movement of a paralyzed eye in terms of the relation between retinal stimulation and the efferent copy of the motor command to move the eye³³.

Thus, different sensorimotor theories emphasise the different levels of sensorimotor contingencies to varying degrees. Noë and O'Regan focus chiefly on contingencies between subpersonal stimulation and exploratory movement, whilst Noë (2004) focuses on those obtaining between personal-level perception and movement, whilst also appealing to contingencies at other levels in various places in the book. Hurley's view is difficult to categorise in these terms – she argues that the personal-level categories of perception and action have a shared underpinning in the dynamic, looping nexus of perceptual inputs and

³¹ See (e.g.) Weiskrantz (1998) on gaze-redirection in blindsight patients.

³² Though at other points in the book different levels of sensorimotor contingency are emphasised, as when he appeals to the role of patterns of variation in retinal stimulation due to inhomogeneities in the distribution of rods and cones in his account of colour perception (p.129-130).

³³ Though Hurley usefully uses 'perception' and 'action' to refer only to personal-level phenomena, talking instead of 'input' and 'output' at the subpersonal level.

motor outputs at the subpersonal level. At the subpersonal level, inputs and outputs cannot be cleanly distinguished, and the fact that the contents of perception and action each supervene on this single subpersonal system means that they are both co-enabled and co-defined. Such a framework seems able to appeal to sensorimotor contingencies at any of the levels surveyed above³⁴. For our purposes, we can define a sensorimotor theory as any view which attempts to explain the content and character of perception via a non-instrumental appeal to any or all of these types of dependencies between perception and action.

2.2 The Dual Visual Systems Hypothesis

The Dual Visual Systems hypothesis (DVS) (Milner and Goodale (1995), Jacob and Jeannerod (2003), Goodale and Milner (2004), Jeannerod and Jacob (2006)) identifies two distinct strands of visual processing in the human brain. The ventral stream runs from the primary visual cortex, through V4, to the temporal lobe, and appears to be implicated in object-recognition and classification, and linked to working memory and capacities for visual imagery. The phylogenetically older dorsal stream runs from the primary visual cortex, through MT, to the posterior parietal cortex, and is implicated in our fine-grained motor interactions with the world, such as controlling reaching movements towards objects and calibrating appropriate grip-sizes with which to grasp them. Initial evidence for the existence and possible functions of the two processing streams came from lesion studies on primate vision (with antecedents in work on frogs (Ingle (1973)) and hamsters (Schneider (1969)) by Mishkin and collaborators (Ungerleider and Mishkin (1982)), and has subsequently been well confirmed as a neuroanatomical hypothesis by various single cell (e.g. Taira et al (1990)) and neuroimaging (e.g. James et al (2003)) studies³⁵.

However, neuroanatomy is not our concern here - we are interested in what DVS suggests about the functional role of conscious vision, and how this might diverge from the role of motion-guiding vision. Perhaps it is useful to begin our survey of DVS by considering what our visual contact with the world allows us to do. It is a commonplace that vision allows us to, among other things, recognise, recall and reason about perceived objects, to contrast

³⁴ See Hurley (1998), esp. ch.10, for detailed discussion of her 'two-level interdependence view'. I discuss the relation between Hurley's position and the action space view in chapter five, but focus chiefly on Noë's (2004) position until then, since it affords the best contrast with the action-space view I develop.

³⁵ See Jacob and Jeannerod (2003) ch.2 for a useful survey of the neuroanatomical findings in primates.

them with other perceived and unperceived objects, to see how they might help or hinder the satisfaction of various of our goals and projects, and to fluently act upon them by, for example, grasping them with just the right hand-position and finger-separation for their shape. However, note that the very last capacity seems to demand very different computational resources from the others on the list. To identify a perceived object as of a particular type, to remember that object when it is no longer perceived, and to reflect on its actual or possible relations to other objects or to our intentional goals all require us to (in some sense) grasp the fact that it is an objective and enduring object, whose existence and properties are independent of our current perceptual relation to it – its objectivity must be in some way represented, or factored into, our perception of it. On the other hand, to quickly and fluently grasp an object, we need only process information about its current location relative to ourselves (in particular to the bits of our body we wish to use to grasp it), and the properties of it relevant to the way we should grasp it, such as its width and orientation. The availability of such information doesn't require width and orientation to be represented as objective and enduring properties of an external object – just that those features be processed in a way suitable to guide fluent action in the here-and-now. But this in turn seems to require that information regarding the position of the object (or at least of the portions of it to be grasped) relative to the agent be constantly updated as the hand and arm move toward it, that its position and orientation be coded with respect to the reaching arm, wrist and hand in a way that allows each to be positioned and oriented in a way appropriate for grasping the object, and for all these parameters to be computed with sufficient accuracy to support the speed and precision with which we can grab a perceived object. None of these computational demands seem necessary to perceive that an enduring object has a certain property. Thus, these two functions of vision – guiding fluent action on one hand, and enabling recognition, recall and reasoning on the other – seem to require very different computational resources³⁶:

“The former requires a constantly updated, (multiply) egocentrically specified, and exquisitely distance and orientation-sensitive encoding of the visual array. The latter requires the computation of object-constancy (objects do not change their identity every time they move in space) and the recognition of items by category and significance, irrespective of the fine detail of location, viewpoint, and retinal image size. A computationally efficient coding for either task looks to preclude the use of the very same encoding for the other.” (Clark (2007), p.565-6)

³⁶ See Milner and Goodale (1995), p.25-66, for a more detailed discussion of such computational arguments for DVS.

The crux of DVS is that each of the anatomical pathways sketched above specialises in handling only one of these sets of computational challenges, the dorsal stream dealing with the former, and the ventral with the latter.

Strong support for this hypothesis comes from dissociation studies of patients with damage to one of these streams but not the other. Visual agnosics, such as Milner and Goodale's patient DF, have damaged ventral streams but intact dorsal streams. As a result, DF is unable to visually recognise or classify objects, and cannot perceptually discriminate between even simple triangular, circular and square objects, reporting no conscious experience of their shapes. However, whilst denying experience of these shapes, she can fluently pick them up using appropriately-scaled grips on the basis of her visual exposure to them. In a striking demonstration of DF's spared and impaired abilities, Milner and Goodale found that she was virtually at chance when asked to orient a hand-held card to match the orientation of a visually presented slot (whose orientation is varied over the course of the trials), but could post the card through the slot with almost the same degree of accuracy as a normally-sighted control subject³⁷. Also noteworthy is the fact that DF had to be prompted by the experimenters in order to successfully post the card – she denied that she could see the slot, and was initially surprised when her posting behaviours were successful³⁸. And, just as DF was unable to orient the card successfully whilst refraining from actually posting it, she was unable to orient it successfully when a delay was introduced between the presentation of the slot and the instruction to post the card. The hypothesis is that each of these last two conditions constitutes 'pantomimed' action, guided by ventrally-mediated visual imagery or memory, and thus compromised in DF.

In contrast to visual agnosics, optic ataxics have damaged dorsal, but intact ventral streams. Whilst they can visually recognise and classify objects, their performance of the fluent reaching and grasping movements available to DF and normal subjects is impaired. Jeannerod et al (1994) found that optic ataxic AT could provide an accurate manual estimate of object-size by matching it to the separation between her thumb and index finger, but could not accurately scale the grip between thumb and index finger in order to

³⁷ See Milner and Goodale (1995) for a full account and discussion of these results.

³⁸ Matters are slightly complicated when interpreting the abilities of DF by the fact that she has spared colour and texture vision (thought to be due to spared regions of her ventral stream). Thus she can often guess the identity of an object, and the resulting ways in which she can act on it, by inference from these cues. Goodale and Milner (2004) report that she has also developed an ability to self-prompt her actions by making exploratory movements and seeing how they unfold. As a result of her experience of her condition, and her subjection to numerous psychophysical tests, she has a good grip of the sorts of actions she can and cannot fluently perform, and thus no longer exhibits surprise at her ability to perform them.

grasp the same object. We would thus expect that optic ataxics attempting DF's posting task could consciously perceive, report and remember the orientation of the slot, and consequently succeed in the pantomimed and delayed-posting conditions, whilst showing an impairment in the actual posting of the card through the slot³⁹.

Additional support for DVS's view of the division of labour between dorsal and ventral streams comes from the fact that different aspects of normal human vision are differentially susceptible to certain illusions, with some illusions apparently deceiving only vision devoted to categorising or attributing properties to objects without fooling motion-guiding vision. In the Titchener circles illusion, two circles of the same size are experienced as being different sizes when one circle is surrounded by an annulus of smaller circles and the other by an annulus of bigger circles (the circle surrounded by smaller circles appears to be bigger). Aglioti et al (1995) constructed a version of the illusion with poker chip shaped discs, and demonstrated that whilst subjects have a non-veridical experience of the sizes of the discs, as evinced by their choosing the chip surrounded by smaller discs when asked to pick up the larger of the two central chips, their grip will be accurately scaled to the actual size of the disc, rather than to the illusory size which it is experienced as having. Grip is scaled identically if the subjects are asked to pick up the other, identically-sized but apparently smaller, disc. Thus, it seems that their perception of the relative size of the discs and their resultant choice of which disc should be picked up in order to comply with the experimenter's request are under the influence of the illusion, whilst their reaching behaviour is not. Daprati and Gentilucci (1997) have obtained similar results for the Muller-Lyer illusion – when subjects attempt to pick up either of the central bars in the illusion, their grip is scaled appropriately, despite their erroneous perception of the bar's size, as evinced by drawings or manual estimations⁴⁰.

³⁹ Jacob and Jeannerod (2003, p.186) thus note that AT is aided in the task of pointing to a visually presented light by a delay between the presentation of the light and the pointing. Matters are complicated when contrasting optic ataxics with visual agnosics by the fact that ataxics usually show only slight behavioural impairments, often only relating to stimuli in their peripheral rather than central vision. Rosetti (forthcoming) suggests that this casts doubt on DVS's view of the division of labour between the streams, but the paucity of most impairments might also be explained by the fact that dorsal lesions are usually relatively minor – due to its location in the brain, a lesion that fully compromised the dorsal stream would result in death. It seems reasonable to say that subjects with the most severe lesions approximate most closely the deficits that a strong version of DVS would predict.

⁴⁰ The dissociations Daprati and Gentilucci found between perceptual and grasping conditions were less striking than those in the Titchener circles experiments. This is likely explained by the fact that their experiment involved presenting the bars in the Muller-Lyer illusion singly rather than concurrently – thus perceptual estimation of size was not affected by comparison of two illusory lengths, as it would be in the standard Muller-Lyer illusion. It seems reasonable to suppose that concurrent presentation of the bars would

Conversely, some perceptual cues seem to have a far greater effect on reaching and grasping than on perception and judgement. Haffenden and Goodale (2000) and Haffenden et al. (2001) show that the distance between the central circle and the annulus in the Titchener circles illusion has a far greater effect on grip-scaling than on perceptual estimates of size⁴¹. They hypothesise that this is because the visuomotor system that controls the grasping movements treats the annuli as obstacles to be avoided when reaching, whereas this consideration is irrelevant to the purposes of the ventral processing which causes the perceptual illusion.

The last class of illusions we should consider *do* have an effect on the way perceivers engage with the non-veridically perceived object, but not on the scaling of the grip used to act upon it. In the Ponzo, or 'railway lines' illusion, one or two objects are displayed within a pair of converging lines. The lines give an impression of receding into the distance towards a vanishing point, and objects placed further up between the lines appear larger than identically-sized objects lower down. Brenner and Smeets (1996) asked subjects to pick up discs placed at various points in between the lines, and found that the force used to lift the disc increased the higher up the disc was placed, suggesting that subjects took the disc to be larger, and thus heavier, than it in fact was. Similarly, Ellis et al (1999) found that when a bar is placed at an angle between the Ponzo lines, subjects grip it closer to one of its ends, suggesting that they take it to be heavier at one end than the other, and Jackson and Shaw (2000) found that a bar was grasped with greater force according to how high up the lines it was placed. Grip-scaling remained unaffected in each of these experiments.

These illusions and the studies on lesioned subjects support the idea that there is a distinction, and possibility of dissociation, between the fine control of reaching and grasping behaviour and other aspects of visual perception, and we have good reason to believe that this corresponds to the distinction and possibility of dissociation between the ventral and dorsal streams of processing. What can this tell us about the functional role of conscious visual experience? The first point to note is that ventral stream activity seems especially implicated in visual consciousness. DF's conscious experience of visual form was

increase the illusory size discrepancy, and produce a stronger dissociation between perceptual and grasping conditions.

⁴¹ These studies are two of several designed to resolve the controversy over Aglioti et al's initial results. Those results are controversial since slight effects on grip-scaling are in fact found in such experiments. The upshot of the considerable literature on this topic seems to be that this can be explained by appeal to the visuomotor system treating the annuli as objects to be avoided, and interaction between dorsal and ventral streams – I'll thus assume here that the results do indeed support the version of DVS under discussion. See Jacob and Jeannerod (2003, ch.4) for a review.

lost when her ventral stream was damaged, as evinced by her reports of its absence, her initial surprise at her ability to perform certain types of action, and her inability to use visual information about objects to initiate actions upon them based upon their shape, or to recognise or recall their shape-properties. Likewise, our illusory conscious perceptions of the figures in the various illusions discussed above seem to be due to ventral processing, rather than the dorsal processing which appears to accurately track the attributes about which it processes information⁴². If we grant that we have reason to think that visual consciousness is underpinned chiefly by the ventral stream, what do the above results suggest about the function of visual consciousness?

Firstly, it seems plausible that conscious vision is responsible for making objective and enduring properties of objects apparent to the perceiver. In this connection, Matthen (2005) contrasts 'descriptive vision', devoted to building up a record of the objective properties of visually-presented objects that is accessible to memory and conscious awareness, with motion-guiding vision, devoted to the control of fine details of action on visually-perceived objects. The following aspects of the DVS data support the idea that the ventral stream, and thus conscious vision, is geared towards representing objects in the former way:

- The computational arguments considered at the start of this section suggest that the ventral stream must somehow represent, or take account of, such properties of objects.
- DF is unable to categorise, identify or remember items based on their objective shape, and reports no awareness of that shape, whilst nonetheless being able to display sensitivity to their shape when prompted to grasp them. This suggests that to the extent to which DF's brain processes information about shape, it treats it egocentrically, representing an object's shape only in terms of the movements DF needs to make in order to grasp it, given her current position relative to it.
- By contrast, the ventrally-intact optic ataxic's visual sensitivity to shape allows her to reason about, recognise and recall the consciously-perceived shape (all processes that we suggested above require shape to be represented as an objective, perceiver-independent property of the object), whilst the way in which information about the relations between object's shape and the current position of the perceiver can be put to use in grasping an object is impaired.

⁴² It might be objected here that the supposition that the activity of the dorsal stream does not issue in consciousness is unwarranted – perhaps it issues only in conscious experience that is somehow unavailable to report and judgement. I address this objection in section five.

- The generally accepted explanation of the non-veridical experiences in ventrally-mediated illusions is in terms of the visual system's attempts to maintain size constancy across the visual field, triggered by illusory perspectival cues. In the Titchener circles illusion, the annulus of small circles is treated as composed of far-away objects, and the central circle is taken to be the same distance away. It is thus perceived as larger than the other central circle, which is perceived as nearby due to its being surrounded by an annulus of larger circles. Similarly, the lines in the Ponzo illusion are perceived as being different distances away from the perceiver, and hence the line higher up the Ponzo tracks is perceived as larger. This suggests that our conscious perception of size results from the ventral stream's treating it as an objective property, which an object possesses independently of the particular angle it currently subtends in our visual field.⁴³
- Similar considerations apply to the demonstration by Brenner and Smeets (1996) and others that the position of objects in between the Ponzo lines affects the weight they are perceived as having.

Secondly, it seems plausible that conscious vision is linked to capacities for memory and action-selection in a way that motion-guiding vision is not. In this connection, Clark (2001) proposes the "Principle of Experience-based Selection" on the basis of consideration of the DVS data, according to which "Conscious visual experience presents the world to a subject in a form appropriate for the reason-and-memory based selection of actions" (p.512), as opposed to a form appropriate for the fine-grained control of action, such as the calibration of reaching and grasping movements. Several parts of the DVS data support this claim:

- Our initial computational arguments suggest that ventral stream processing is apt for integration with reasoning and memory, but that the business of fine-grained control requires a different set of informational resources.
- Whilst DF can grip objects accurately, she shows impairments in selecting the appropriate type or target of the grip – for example, she might deploy a precision rather than a power grip when picking up a screwdriver, and pick it up by its tip rather than its handle.
- DF's perceptual sensitivity to the orientation of the slot in the posting task is not such as to allow her to reason or plan that she would have to orient her hand *like so* were she to post the card – we noted above that she fails at this task. Nor is it such as to allow her to post the card through the slot after a delay between seeing the slot and attempting the action, suggesting that the information processed about

⁴³ See Jacob and Jeannerod (2003, ch.4) for discussion of such accounts of the illusions.

the slot's orientation is short-lived and not accessible to memory, or systems that would enable her to plan or reason that the card should be posted with a certain orientation.

- Dorsally damaged but ventrally intact optic ataxics consciously experience the slot's orientation, but show the converse set of impairments – they do better at matching the orientation of their hand to the orientation of the slot, or posting the card after a delay, than they do at the straightforward posting task.
- The illusory experience of the relative sizes of the Titchener circles guides the subject's selection of which disc to pick up in Aglioti et al's (1995) study, but not the scaling of the grip they use.
- In the studies with the Ponzo illusion, the force the subject elects to use when picking up the bar is affected by the illusion, but the scaling of their grip is not.

The DVS data might inspire other suggestions about the nature and function of conscious vision. For instance, Jacob and Jeannerod (2003, e.g. p.112-3, p.193-8) propose that the DVS data suggests that conscious vision is essentially comparative and contrastive, arguing that the ventral stream always represents properties of objects in relation to other real or possible objects (in an 'allocentric' frame of reference), whereas the dorsal stream represents properties of objects only in relation to the fine details of certain possible actions upon them by the agent (an 'egocentric' frame of reference). I won't pursue this, or other suggestions, further here – it seems to me that the conclusions that the DVS data warrants most clearly are captured under our two headings above. So, with an overview of the DVS data and its likely upshots in hand, let us now consider its relevance to enactive theories of perception.

2.3 DVS and Sensorimotor Theory

Several commentators (e.g. Block (2005), Clark (2006a, 2006b), Jacob (2006)) have claimed that there is a tension between Noë's sensorimotor theory of perception and the DVS results. And at first glance, the relevance of these results to the sensorimotor view seems fairly direct. DVS apparently demonstrates at least a functional dissociation between conscious visual perception and our visually-guided motor interactions with the world. By arguing for the constitutive role of action and sensorimotor engagement in visual perception, it seems that the sensorimotor view is emphasising the functions of vision thought to be dealt with by the dorsal stream. However, a natural interpretation of DVS suggests that these functions are largely independent from our conscious, reportable,

memory-informing perception. Thus, the plausibility of the sensorimotor theory as an account of the basis of visual consciousness is under apparent threat.

Hence, Block notes the fact that colour and size perception are significantly worse in our peripheral visual field than in its centre, but that our capacity to reach and grasp objects on the visual periphery is largely unimpaired, and concludes that “The upshot is that if the activity guided by sensorimotor knowledge with which the enactive approach includes visually guided action, it *simply does not reflect the phenomenology of conscious vision.*” (Block (2005), p.10, his italics). Jacob (2006, p.9-10) makes a similar point, claiming that the sensorimotor view appears to be committed to characterising visual experience as tied to the control of action, a tie of which DVS gives us reason to be suspicious. And Clark uses the DVS results to argue for a “sensorimotor summarising” view of conscious perception, where the function of conscious vision is geared towards reasoning, planning and the selection of actions. Once such action-plans have been formed, the job of carrying them out in fine detail is delegated to dorsal stream processing that does not issue in conscious awareness. He concludes that dealings with detailed sensorimotor knowledge are the preserve of the dorsal stream, and thus of limited importance in explaining perceptual consciousness.

Hurley and Noë (2007) give a brief response to such worries about the possible dissociation between conscious and motion-guiding vision, focussing on the Titchener circles, and other ‘action resists illusion’ results surveyed in the previous section. They agree that such findings seem to constitute cases where the subject’s implicit sensorimotor knowledge and expectancies, but not their conscious experiences, are tuned to the actual size of (for example) the Titchener discs, as evinced by contrasting their intact abilities to act accurately upon them with their non-veridical phenomenology and judgements of relative size. However, they argue that to reject their sensorimotor view on this basis is premature – the results show only that a narrow range of sensorimotor knowledge and skills are left intact by the illusions (such as the practical knowledge of how to scale one’s grip appropriately to the size of a disc), whereas the range of sensorimotor knowledge and expectancies to which they wish to appeal in explaining the content and character experience is far broader than this. Thus, they claim that:

“Subjects in the experiment of Aglioti et al can use what they see to guide certain specific movements correctly [...] But that doesn’t show that the broad set of sensorimotor powers and expectancies associated with the illusory look of

Titchener circles is significantly different from the set associated with similar but veridical looks in a variety of different contexts.” (Hurley and Noë (2007), p.18)

Moreover, they object that even the narrow range of illusion-resistant abilities and expectancies may be “highly specific to certain tasks or contexts” (Ibid, p.19) such as carefully prepared laboratory set-ups, and that we should resist drawing any general morals from such cases as a result.

As it stands, this is not a satisfactory way of reconciling DVS and the sensorimotor view⁴⁴. Firstly, it leaves the sensorimotor view in need of an account of why some sets of sensorimotor knowledge and skills should be more or less relevant to the character of consciousness than others, which H&N do not provide. Even if some property separating sensorimotor knowledge that is pertinent to experience from sensorimotor knowledge that is not could be plausibly specified, the sensorimotor view would risk losing its distinctiveness by explaining experience in terms of that property, rather than in terms of sensorimotor features⁴⁵. More importantly, H&N fail to note that the action-resists-illusion results are part of the larger body of DVS data, which points to a pervasive dissociation between two types of vision, going far beyond the specific contexts in which those results are obtained. The real challenge to the sensorimotor view thus arises from viewing these results in the context of DVS; can the sensorimotor view accommodate the distinction between conscious and motion-guiding vision that the full spectrum of DVS data suggests? H&N’s response here does not address this crucial question.

We can make better progress in considering the compatibility of the sensorimotor view with DVS by considering some brief remarks from Noë (2004). There he suggests that the theories are not in conflict so long as we distinguish between deployment, or possession, of

⁴⁴ I don’t want to deny that H&N’s comments here are correct, as far as they go. As we shall see, the account I develop suggests that the content and character of our experience is to be explained in terms of a broader range of abilities than reaching and grasping behaviours, in keeping with their suggestion. However, I will argue that the correct way of cashing this claim out supports the action-space view, and gives us reason to reject the sensorimotor theory.

⁴⁵ Again, it will become apparent over subsequent chapters that this is what the sensorimotor theory must do. H&N’s only gesture at what such a property might be is a brief suggestion, made whilst developing their point about the restriction of the experimental results to certain contexts, that the ‘implicit’ rather than ‘explicit’ use of information (see Hurley (1998), ch.4, for an account of this distinction) in these cases may contribute to an explanation of why the information that guides grasping does not seem to issue in experience. As we shall see in chapter 5, this suggestion is in keeping with my eventual account, and thus in fact favours the action-space account over the sensorimotor view.

sensorimotor skills, and the fine-grained control of action which the dorsal stream is thought to handle⁴⁶. Hence he claims that:

“[Optic ataxia] does not undercut the enactive view, for from the fact that a patient suffers optic ataxia, it doesn’t follow that he or she lacks the relevant sensorimotor knowledge. What would undercut the enactive approach would be the existence of perception in the absence of the bodily skills and sensorimotor knowledge which, on the enactive view, are constitutive of the ability to perceive.” (Noë (2004), p.12)

Similarly, he says later that:

“[DVS] is, at best, orthogonal to the basic claims of the enactive approach. The enactive approach is not committed to the idea that vision is for the guidance of action, so neither the fact that some visual processing *is* for the guidance of action, nor the fact that some visual processing *is not*, has any direct bearing on the enactive approach. From the standpoint of the enactive approach, all perceptual representation, whether the result of dorsal or ventral activity, depends on the perceiver’s deployment of sensorimotor skills.” (Ibid. p.19)

The idea, then, is that knowledge and expectancies concerning the way appearance changes with movement are being appealed to as constitutive of perceptual experience, not the way in which perception so-constituted is then put to use in the control of action⁴⁷. However, the possibility is still open to Noë of holding that possession or exercise of some type of sensorimotor knowledge is constitutive of perceptual consciousness (and perceptual representations in the ventral stream) without holding that this need take a form suitable for fine-grained control of action.

⁴⁶ Noë’s official position seems to be that the *deployment* of sensorimotor knowledge is constitutive of the content and character of perception, though occasionally he writes as if mere *possession* of such knowledge were sufficient. I won’t attempt to distinguish or adjudicate between these two options here.

⁴⁷ Note that dorsal perceptual representation could not depend on sensorimotor knowledge in the preferred sense of Noë (2004), where the perceiver implicitly understands how (phenomenological) appearances will change with movement, since visual form agnosics (in conjunction with other aspects of the DVS data) seem to demonstrate that dorsally-mediated vision takes place without the subject enjoying any visual experiences. It might be that a different type of sensorimotor account could be given for the dorsal pathway, perhaps giving ‘knowledge’ and ‘expectancies’ some subpersonal construal.

Is this position a stable one? Jacob (2006) notes that Noë makes this move, but questions whether all aspects of Noë's theory are consistent with the eschewal of appeal to visually-guided action. The passage Jacob cites (Ibid. p.9) is Noë's discussion of the role of egocentric behavioural space in the localisation of objects and our experience of the way they look. For example, Noë claims:

“The plate looks elliptical to me because, to indicate its shape, I can (and indeed, in some sense, *must*) move my hand in a characteristic manner. That is, to experience a thing as elliptical is precisely to experience it as occupying a particular kind of region in one's egocentric, sensorimotor space.” (Noë (2004), p.89)

Presumably Jacob's idea here is that in tying our visual experience of a shape to possible bodily indications of that shape, Noë is committed to constitutive ties between motion-guidance and visual experience that run counter to the DVS data. However, it is not clear that this passage, or this type of appeal, provides Jacob with the ammunition he needs. Noë is here advocating an account of perception of spatial properties which ties it in some way to possible behaviours of the subject, but this does not commit him to the counterintuitive claim that the ensuing visual representations are *for* these behaviours. At most, it commits him to the claim that an appeal to such behaviours plays an essential role in a proper specification of the content of spatial perception – and it is not clear why this claim should be threatened by DVS. Moreover, it is not clear that the behaviours to which Noë appeals here are the fine-grained motor responses which the dorsal stream coordinates. Moving one's hand through the air to mime the contours of a visually-perceived shape in the way Noë describes looks like just the sort of pantomimed action that the dorsally intact DF cannot do, but that ventrally intact but dorsally compromised optic ataxics can.

In the other passage that Jacob cites in support of his claim that Noë's view is committed to an action-guiding role for visual experience, Noë notes that “there is no question that vision can and does guide movement, but it is not the business of the sensorimotor view to emphasise this humdrum fact” (Ibid. p.90). Jacob responds that “The evidence for the two-visual systems hypothesis is precisely what would lead one to doubt at least the strong assertion that the function of perceptual experience is to guide movement” (Jacob (2006), p.9). However, this does not seem to be the assertion Noë makes in the quoted passage, which looks compatible both with the view that only dorsally-mediated vision guides movement, or a more sophisticated view which holds that conscious vision controls

movement only in the sense of selecting and delegating actions to be controlled by the dorsal stream.

So, Jacob fails to make a case that Noë is committed to making claims in direct tension with DVS. The key point is that in none of the passages Jacob cites (nor, as far as I can tell, throughout Noë (2004)) does Noë appeal to the implicit knowledge that concerns the relationship between sensation and the kinds of fine-grained reaching and grasping movements which are under dorsal control, knowledge that DVS suggests is largely independent of our conscious visual experience. Rather, the types of “skilled bodily activity” to which he appeals to are capacities for intentional exploratory head, body and eye movements. For Noë, conscious perception is borne out of a perceiver’s understanding of the relationships between these types of actions and changes in perceptual experience. Since DVS gives us no reason to think that these sorts of actions or understanding are dorsally governed, Noë cannot be straightforwardly convicted of the most intuitive DVS-inspired objection to his view – that it is in tension with DVS by dint of emphasising capacities that look to be the province of the dorsal stream in its account of perception consciousness.

It seems that analogous remarks apply to each of the variants of the sensorimotor view we considered in section 1. Noë and O’Regan’s (2001) view emphasising implicit knowledge of the relationship between patterns of subpersonal stimulation and exploratory movement looks compatible with DVS so long as they do not make an appeal to the types of actions thought to be under dorsal control and subpersonal predictions about their sensory consequences in their account of visual experience. The high-level sensorimotor contingencies appealed to in Myin, O’Regan and Noë’s (2002) treatment of certain cases look orthogonal to any theory of fine-grained motor control. And Hurley’s (1998) appropriation of Gallistel’s paralyzed eye case for the sensorimotor theory looks compatible with DVS even if we assume that gaze-redirection can be a dorsally-mediated function, since here the appeal is to the subpersonal consequences of an *intentional* action (saccading *intentionally* to the left, rather than doing so as a result of automatic dorsally mediated gaze redirection) in explaining perceptual experience. More generally, Hurley’s contention that the subpersonal inseparability of perceptual inputs and motor outputs results in the co-constitution of perception and action at the personal level seems perfectly in keeping with the DVS results, so long as she does not claim that the fine details of certain forms of action-guidance are reflected in conscious experience. DVS is not straightforwardly incompatible with sensorimotor theory.

Block (2005) and Clark (2006) lodge more subtle DVS-based objections to sensorimotor theory. Like Jacob, Block is skeptical about whether Noë is consistent in his denial that the type of visual guidance of action thought to be the function of the dorsal stream is one of the bodily skills to which the sensorimotor view appeals, though he does not attempt to cite any passages in support of this. Instead, Block makes a number of other DVS-based objections, the two most significant of which we will deal with briefly.

Firstly, he notes that “presumably, the dorsal stream involves at least implicit expectations for how objects in the environment shift with visually guided movements”, and thus concludes that Noë is committed to claiming “that the (in fact unconscious) dorsal states are conscious” (Block (2005), p.11). But this does not follow. Firstly, Noë’s (2004) view focuses on expectations concerning the experiential consequences of movement, rather than any ‘implicit expectations’ coded for in the dorsal stream. Secondly, even applied to a sensorimotor view that did focus on sensorimotor ‘expectations’, subpersonally construed, it would follow only that the view needed a supplementary story about why some such expectations but not others were relevant to consciousness, not that the operations of the dorsal stream must have an upshot in perceptual consciousness. Whilst the provision of such a story may be difficult, Block gives us no reason to suppose that it could not be given⁴⁸.

Secondly, Block notes that if Noë’s view emphasises not visually-guided action but sensorimotor expectations (as it indeed does), then Noë will be forced to leave behind much of the radical externalism of his view – expectations can presumably be instantiated in the absence of an acting body, so Noë’s vision of perception as essentially active and embodied would not be supported by his theory of perception. Two points in response – firstly, Block thinks Noë is forced to emphasise sensorimotor expectations since a sensorimotor view which appealed only to *dispositions* towards visually-guided actions is also rendered off-limits by the DVS data. Block makes this point in the context of analysing Noë’s denial that optic ataxia – where visually-guided action is compromised but visual experience is not – undercuts the sensorimotor view, and his subsequent claim that:

‘What would undercut the enactive approach would be the existence of perception in the absence of the bodily skills and sensorimotor knowledge which, on the enactive view, are constitutive of the ability to perceive.’ (Noë (2004), p.12)

⁴⁸ Though recall our point from earlier this section – such a supplementary story risks shifting the explanatory focus in the sensorimotor account of experience from sensorimotor knowledge to whatever property is appealed to in distinguishing between those sets of sensorimotor expectations that are relevant to the content and character of experience from those that are not.

Block thinks that the sensorimotor view cannot cash out an appeal to skilled behaviour in its account of perception in dispositional terms, since it seems plausible that optic ataxia compromises dispositions to skilled behaviour as well as the ability to execute such behaviour⁴⁹. However, as we have just noted, so long as the skills or dispositions appealed to by an enactive view do not include the fine details of grip-scaling and other dorsally mediated operations, we have no reason to think that the view is in tension with DVS - the dispositions putatively lacked by the ataxic are relevant to an enactive view only if that view emphasises dispositions to dorsally-governed action in its account of perception. DVS thus constrains, but does not rule out, a dispositional theory of perception.

Secondly, even if we grant Block that the most plausible form of sensorimotor theory is one in terms of sensorimotor expectations that could be realised in the absence of a body, it is debatable whether this threatens the non-instrumental nature of its appeal to action in its account of perception. Block disagrees, since he thinks questions about what is constitutive of any given perceptual experience should be settled by specifying the minimal metaphysically sufficient supervenience base for that experience (Block (2005), p.5). If perceptual experience depends on sensorimotor expectations, and sensorimotor expectations can be realised by neural activity alone, then Block thinks that skilled activity is only instrumentally important to perceptual experience - important insofar as it causally effects the neural activity that forms the minimal supervenience base of experience. But an enactivist might disagree with Block's criterion for constitutivity and non-instrumentality; why suppose that such questions should be settled by considerations about supervenience, rather than (for example) considerations about the factors which must be ineliminably appealed to in our best explanation of the phenomena? This option is pursued by Hurley (forthcoming) and Noë (2007). Hurley notes that if our primary concern is the explanation and understanding of perceptual experience, the fact that neurally identical brains will enjoy identical experience regardless of their different embodiment and environment is of no great relevance - it tells us nothing about why the resultant experience should be of, for example, vermilion rather than violet. Since we have, at present, little idea of how such explanatory gap questions should be answered, we should not discount the possibility that an explanation of the content and character of the experience must make reference to extra-neural factors, such as the nature of the sensorimotor dynamics and environmental backdrop in which the neural state is embedded. Indeed, argues Hurley,

⁴⁹ Block's reason for thinking this seems to be that optic ataxia occurs as a result of a brain lesion, and thus 'affects the dispositions themselves, not just their manifestation' (Block (2005), p.12). This seems questionable, but I won't pursue the point here.

the plausible diagnoses that the sensorimotor theory affords of various empirical results speak in favour of this possibility⁵⁰. Even if the contentious point that figuring in a minimal supervenience base is criterial of non-instrumentality and constitutivity is granted, Hurley also notes that the ‘internalist’ claim that this supervenience base will include neural properties alone has not been demonstrated. Noting, as Block does, that some account of perception does not rule out internalism shows only that the internalist/externalist debate has not been settled *a priori*. To assume that this constitutes a victory for internalism is to make what Hurley labels “the ‘causal/constitutive error’ error”,

“of objecting that externalist explanations give a constitutive role to external factors that are ‘merely causal’ while assuming without independent argument or criteria that the causal/constitutive distinction coincides with some external/internal boundary” (Hurley (forthcoming), p.5)

Both Block and the Noë can agree that sensorimotor dynamics causally influence neural activity. But Block thinks that we should conclude from this that experience is not essentially active. Hurley’s point is that this begs the question against the enactivist, since the tacit assumption that our default supposition should be that all essential properties of an experience are neurally enabled has not been defended, and that is precisely the point at issue between Block and the enactivist⁵¹.

Clark (2006a) brings us closer to the real challenge posed by DVS for the sensorimotor theory. He draws similar conclusions from the DVS data to those summarised at the end of the last section, focussing on the suggestion that the enabling of various forms of classification, reasoning and planning seem especially relevant to conscious experience, and concludes that:

⁵⁰ Such as the TVSS phenomena discussed in 1.5 and Kohler’s goggle-adaptation results, to be discussed below and in the next chapter. See Hurley and Noë (2003) and Hurley (1998, ch.9) for more discussion of empirical work that a sensorimotor perspective seems to accommodate well.

⁵¹ I have taken some time over these arguments, since they apply to enactive theories in general, and thus to the action-space view as well as the sensorimotor theory. I remain neutral on the internalism/externalism debate for the remainder of this thesis. However, I do hold that the action-space account’s appeal to action is non-instrumental as billed, insofar as even if the minimal supervenience base of an experience *is* neural, we need to know the space of actions the agent takes to be currently enabled and how this relates to the neural properties of the agent’s brain in order to correctly specify the neural properties that are essential to the experience. The case for this claim is made in chapter 6.

“If this is correct, then one of the most striking implications of the full sensorimotor view may be called into question. This is the claim that all differences in embodiment (insofar as they impact sensorimotor contingencies) must thereby make *some* difference to qualitative experience...This claim will turn out to be false if what structures experience is (not the full suite of sensorimotor details but) a kind of coarse summary whose main concern is with the selection of targets and of action types.” (Clark (2006a), p.7-8)

Here, Clark is referring to passages where Noë claims that the fine details of embodiment make a crucial contribution to the character of experience, since many such details will subtly affect the nature of the relationship between perceptual sensation and movement. For example, Noë and O’Regan state:

“For two systems to have the same knowledge of sensorimotor contingencies *all the way down* they will have to have bodies that are identical *all the way down* (at least in relevant respects). For only bodies that are alike in low-level detail can be functionally alike in the relevant ways” (Noë and O’Regan (2001), p.1015)

And Noë claims that

“Creatures with bodies like ours would have systems that are visual in the way ours are. Indeed, only such systems can participate in the identical range of sensorimotor interactions that we participate in.” (Noë (2004), p.159)

I endorse Clark’s interpretation of the DVS results, but disagree that its primary relevance to the sensorimotor view is the threat that it poses to an emphasis on the fine details of embodiment⁵². This is because such an emphasis seems to be a negotiable but counterintuitive and apparently inconsistent aspect of the sensorimotor view, that it would do well to drop. It is negotiable since it is not entailed by the essential feature of the sensorimotor view as I have sketched it – that a non-instrumental appeal to the perceptual consequences of movement is required for a satisfactory theory of perception. The ‘perceptual consequences’ to which a sensorimotor theory appeals might be at any one of various functional levels – retinal, cortical or experiential, for example – and we have been given no compelling reason to think that the lowest such levels are of importance to an

⁵² However, in his (2006b) Clark treats the DVS-based and ‘sensorimotor chauvinism’ objections separately. The arguments of this chapter are fully compatible with his views as expressed there, which I endorse.

account of perception. It is counterintuitive since, as Clark and Toribio (2001) note, such a close tie between perceptual experience and the fine details of embodiment results in a “sensorimotor chauvinism”, according to which creatures identical but for marginally different saccade-speeds, or retinal forms cannot enjoy identical experiences. And it is apparently inconsistent with various points used to elucidate the sensorimotor theory. For example, Hurley (1998), Hurley and Noë (2003) and Noë (2004) cite Kohler’s visual adaptation results (some of which will be discussed in more detail in the next chapter) in support of the sensorimotor view. In one set of Kohler’s experiments, subjects wore goggles which inverted the visual signal with respect to left and right, so objects on the left are seen through the goggles as on the right, and vice-versa. However, after several days wearing the goggles, the vision of the subjects adapts, and they eventually report visually experiencing everything as in its proper place. The sensorimotor explanation of the adaptation is that there is some higher-order similarity between the sets of relationships between perception and movement that obtain before and after donning the goggles that explains the eventual sameness of the way the world looks pre-goggles and post-adaptation. Adaptation, then, is a matter of the perceiver somehow latching onto this higher-order similarity in the face of what is effectively an alteration of their visual physiology. However, this treatment of the case suggests that the way things look can be determined by sensorimotor contingencies that obtain at a level of abstraction removed from the precise details of embodiment, and that apparently very different forms of embodiment (for example visual systems replicating the effects of any of the various sorts of distorting goggles to which subjects can adapt) are compatible with things looking this way. Lastly, an emphasis on the details of embodiment seems out of place in a sensorimotor theory such as Noë (2004)’s, which emphasises the relations between changes in appearance and possible movements – the details of embodiment are only relevant to such a theory insofar as they affect the way things appear, and we have been given no reason to suppose that *all* details in embodiment should do this. So, whilst Clark is correct to note that DVS gives ample grounds for our suspicion of sensorimotor chauvinism, it was a feature we had independent reasons to question, and which sensorimotor theorists can consistently reject⁵³.

I have argued in this section that Jacob, Block and Clark fail to demonstrate the incompatibility of DVS and sensorimotor theory. Sensorimotor theories explain

⁵³ In fairness to Clark (2006a), he is criticising sensorimotor theory as it appears in Noë’s work, in which there are several references to the importance of the details of embodiment, rather than the less-committal form of sensorimotor theory (characterised as a certain view of the nature of the dependence between perception and action) whose compatibility with DVS I am arguing for here.

perceptual experience by appeal to the perceiver's implicit knowledge of sets of dependencies between appearance and skilful movement, but need not claim that either the type of movements or appearance/movement relations appealed to are such as to suggest they are mediated by the dorsal stream – the claim which would be in tension with DVS. The DVS data suggests that a key part of the dorsal stream's role is to enable fluent engagement with objects by processing sensorimotor information, such as coding for the fine details of the action appropriate to the agent's current visual input intentions, but sensorimotor theories are not committed to appealing to these particular sensorimotor relations. Block notes correctly that sensorimotor theorists require a supplementary story about why the implicit sensorimotor 'knowledge' at work in the dorsal stream apparently fails to issue in conscious experience, but does not show that such a story cannot be given, and his arguments that DVS forces the sensorimotor theorist to abandon their non-instrumental appeal to the importance of action are inconclusive at best. And whilst Clark correctly notes that DVS seems in tension with an emphasis on the fine details of embodiment, such an emphasis is an inessential feature of the sensorimotor view as I have sketched it. However, if we grant that sensorimotor theory does not conflict with the letter of the DVS view, might there nonetheless be reason to think that it conflicts with its spirit?

2.4 DVS and the Action-Space Approach

I think there is a case against sensorimotor theory in the offing from the sorts of points we have been considering – most closely tied to Clark's remarks about what the DVS data suggests about the functional role of conscious vision. Recall the two general morals suggested by DVS about conscious vision from section two. Firstly, we noted that ventral stream activity (and by extension, we supposed, conscious vision) seems geared especially to enabling perception of objective and enduring properties of the environment, allowing for objects and properties to be identified, classified and compared in ways that abstract from certain details of the perceiver's current relations to them, in contrast to the egocentric representations of properties the dorsal stream seems to employ. Secondly, we noted that conscious vision seems linked to capacities for reasoning, remembering and action-planning in a way that the motion-guiding vision enabled by the dorsal stream is not. What, if anything, do these morals suggest about how we should adjudicate the debate between the sensorimotor and action-space theorists about the correct way to appeal to action in an account of perception?

Recall that the sensorimotor theorist wishes to emphasise the agent's understanding of the perceptual consequences of movement, whereas the action-space theorist emphasises understanding of the practical consequences of perception. Now, note that in stating the morals we drew from the DVS data, we needed to make no reference to an implicit grasp of the perceptual results of movement – nor does an agent's having a grasp of the sensory consequences of movement seem to follow from their having the capacities we have identified as distinctive of the ventral stream. However, we did appeal to the ability to use perceptual information in practical reasoning – and it seems to follow that an agent with the distinctive ventral capacities will be able to act on her environment in a way informed by the perceived properties of objects and features, and reason about the ways the disclosures of her current perceptual sensitivity might afford the satisfaction of various of her goals and projects. This account of the capacities enabled by the ventral stream looks much like our sketch of the action-space approach from 1.5. At first blush, then, it seems that the DVS results have a stronger affinity with the action-space approach than with the sensorimotor view.

In support of this claim, consider again what DF's perceptual sensitivity to visual form does and does not allow her to do. DF's brain clearly processes information about the width and orientation of objects since she can, when prompted, act in ways sensitive to these properties with fluency and precision comparable to that of normally-sighted perceivers. However, it seems that this information does not figure in DF's conscious experience (as suggested by her sincere reports, and her initial surprise at her residual visuomotor abilities), nor is it available for integration with her capacities to plan, reason and recall based on her current perceptual situation (as suggested by her requiring an external prompt before putting her spared visuomotor abilities to use, and her inability to report, remember or reason about (for example) the orientation of the slot in the posting task). The action-space account suggests that these facts tell us something about conscious vision – that part of what it is to be conscious of some content (that a slot has a certain orientation, for example) is for that content to be automatically available for integration with capacities for practical reasoning⁵⁴. DF, on such an account, fails to be conscious of the slot's orientation since her perceptual sensitivity to that orientation fails this requirement – in order for orientation-information to be used in her behaviour she requires an external prompt⁵⁵. Similarly, the action-space theorist explains the dissociation between grip-

⁵⁴ This gloss on DF's condition is cashed out in much greater detail in chapter 5, especially 5.4.

⁵⁵ As we have already noted (footnote 9), matters are complicated by DF's residual perception of colour and texture, and the fact that she now has considerable knowledge of the nature and extent of her spared visuomotor abilities. Thus she can now perform many

scaling and conscious perception demonstrated in Aglioti et al's experiments by virtue of the fact that the detailed width information used to scale the subject's grip is not apt to inform their practical reasoning, unlike the non-veridical information about the relative sizes of the discs processed by the ventral stream, as evinced by the fact that the non-veridical information guides the subject's selection of which disc to pick up, but not the scaling of the grip they use to do so.

It seems, then, that the action-space view's contention that perceptual experience must be understood in terms of the perceiver's grasp of the practical consequences of the current deliverances of their perceptual sensitivity accords well with the DVS data. Though we saw in the last section that the sensorimotor theory appears consistent with that data, nothing in that theory seems apt to explain the dissociations between fluent behaviour and conscious experience that DVS highlights as readily as the action-space account's view of the relationship between action and experience. So, the key claim I want to make here is that the DVS data presents a real challenge to the sensorimotor theory *only when viewed in the context of the action-space approach*. If we hold that conscious experience is to be explained in terms of a perceiver's knowledge of the practical consequences of her perceptual sensitivity, we can predict and explain the dissociations suggested by DVS as above – DF is not conscious of visual form in spite of her spared perceptual sensitivity to aspects of that form, since she, unlike standard perceivers, lacks an automatic grasp of the practical consequences of that sensitivity. Standard perceivers have an illusory experience of the relative sizes of the Titchener circles rather than an experience that reflects the veridical information they use to guide the fine details of their actions upon them since the latter information is not available to guide their practical reasoning (such as their choice of which disc to pick up), unlike the information which distorts estimates of size due to illusions of perspective and size-constancy. The DVS results support the action space view since, on their most natural interpretation, those results conform to what we would expect if the action-space account were true. In contrast, whilst the sensorimotor theorist's view of the relationship between action and perception is compatible with those results, no obvious explanation of them is suggested by that view.

fluent actions without prompting. This makes it appear, contrary to our gloss here, that her perceptual sensitivity to the properties of objects and scenes relevant to her fluent activity *is* such that it can be integrated with her ongoing practical reasoning. The action-space account claims that all such cases are to be explained by DF's employing an inferential strategy – either inferring from colour or texture cues that some action is currently appropriate, or performing speculative exploratory actions and seeing how they turn out. The integration of width and orientation information with her ongoing activity is thus, I claim, inferential and non-automatic, in contrast with normally conscious perceivers.

According to the action-space theorist, then, the feature of the sensorimotor view threatened by DVS is the key claim that an understanding the perceptual consequences of movement is of *non-instrumental* importance for an account of perceptual experience. Whilst we saw in the previous section that sensorimotor theory can remain compatible with DVS, for it to *explain* the DVS data it would have to demonstrate that an understanding of the perceptual consequences of movement was required in order to manifest the capacities which seem distinctively enabled by the ventral stream, and are thus associated with conscious experience. However, we have been given no reason to suppose that this is so. Moreover, even if the ventrally enabled capacities of conscious perceivers are in fact grounded in some species of knowledge of sensorimotor dynamics, this does not negate the advantage the action-space view has over the sensorimotor theory in accommodating the DVS data. For it seems that the functions of the ventral stream suggested by DVS could be realised in a variety of different ways – perhaps underpinned by implicit knowledge of sensorimotor dynamics, perhaps by a suitably tuned connectionist network, or perhaps by some suitable modular GOFAI architecture. If we take the morals of DVS seriously as hypotheses about the functional role of conscious vision then each of these suggestions, even if true, would be of importance only as a story about the way the capacities crucial to conscious vision are realised, or enabled. Thus, even if the correct explanation of our possession of those capacities turns out to appeal to knowledge of sensorimotor dynamics, that appeal plays only an *instrumental* role in our account of conscious vision – the sensorimotor dynamics are only relevant insofar as they serve to enable a particular set of capacities. It is those capacities themselves, however enabled, that would do the real work in our account of conscious vision.

DVS suggests that the capacities we should focus on are those enabled by the ventral stream, and I have argued in this section that such capacities are more naturally understood in terms of a perceiver's grasp of the practical consequences of the disclosures of her perceptual sensitivity, rather than of the perceptual consequences of her movement⁵⁶. According to the action-space account, conscious perception consists in a perceiver's manifesting a grasp of the space of actions that the disclosures of her current perceptual sensitivity makes appropriate – a grasp that enables the possibility of those actions to be

⁵⁶ Note that this argument for the mere instrumental or enabling role of sensorimotor dynamics in perceptual experience does not succumb to the criticisms from Hurley and Noë we considered in the previous section. Our relegating sensorimotor dynamics to such a role is based on empirically motivated arguments rather than considerations about minimally sufficient supervenience bases, and we have not traded on the implicit premise that the neural enabling of experience should be our default assumption.

factored in to her ongoing practical reasoning. Over subsequent chapters the nature of this proposal will be clarified, and we will look at independent reasons for the view. But note for now that such a view allows us to make good sense of the dissociations revealed by DVS, and that if we have independent reason for holding such a view, DVS counts as empirical support in its favour. Neither of these points hold for the sensorimotor view – this is why, in the debate between sensorimotor and action-space approaches, DVS favours the action-space camp.

2.5 DVS and Consciousness

The arguments of this chapter have presumed, in line with most interpretations of the DVS literature⁵⁷, that the functions of the ventral stream are of special relevance to visual consciousness, whereas the functions of the dorsal stream are not. Our main reasons for this have been DF's sincere reports that she has no conscious experience of visual form, her initial surprise when confronted with the extent of her spared visuomotor abilities in the face of her apparently impoverished visual experience, and reflection on the way normal perceivers (ourselves included) have illusory experiences of the sizes of figures and objects in certain illusions despite our near-perfect visuomotor abilities with respect to those figures and objects. However, does this give us sufficient reason to suppose that the information guiding visuomotor action in these cases is not reflected in conscious experience? Perhaps, it might be objected, the motion-guiding information is *experienced*, but simply not *noticed*, by the subject. Wallhagen gives this example in support of the possibility of such an interpretation:

“...suppose you see a friend along the street, and your friend has had her hair cut. You might not *notice* this for some time. Nevertheless, you were *aware* of her hair, and the properties in virtue of which it looks different to you (which explains the thought you're likely to have when you *do* notice her haircut: 'I had a sense there was something different about you today')...” (Wallhagen (2007), p.552)

Perhaps, then, DF is aware of the information she uses to guide her visuomotor behaviour, but simply fails to notice it in a way that is reflected in her sincere introspective reports. Certainly the fluency with which she can get around in the world gives us apparent reason

⁵⁷ E.g. Goodale and Milner (2004), Jeannerod and Jacob (2003), Clark (2007). See sections 5-7 of Clark (2007) for his own take on the issues of this section, which is fully compatible with the arguments I offer here.

to suppose she has experience of visual forms – why should her reports to the contrary be given ultimate authority? Likewise, perhaps when we view the Muller-Lyer, Titchener circles, or Ponzo illusions we have both veridical and non-veridical experiences of relative size, with the veridical experience guiding our reaching behaviour and the non-veridical experience guiding our report, recall and practical reasoning.

This is an extremely thorny issue, and I can't hope to settle it here⁵⁸. But note first of all that the introspective reports of DF and subjects perceiving illusions, as well as DF's initial surprise at the dissociations between her reportable visual experience and certain of her visuomotor skills, and our surprise at the dissociation between our reportable experience of the relative sizes of figures in certain illusions, and our visuomotor skills with respect to those figures surely give good *prima facie* reasons to think that the information guiding the details of visuomotor behaviour is not part of our conscious visual experience. Look long and hard at (let's say) the Ponzo illusion. Introspect as thoroughly (or as casually, depending on which you think gives the best insight into your experience) as you wish. Does *anything* in your experience of the lines, theoretical knowledge of the illusion aside, tell you that they are the same length? It seems to me that, so long as I am not straining to ignore the fact that the lines are situated within the Ponzo tracks, there is no sense in which they appear the same length. The objection at hand suggests that I am deceived – that I *do* have a veridical experience of the relative size of the lines, as would be manifested were I to try and grasp them. But what motivation is there to suppose that the information that guides my grip-scaling figures in my experience? Perhaps the fact that it is mildly counter-intuitive to hear from DVS that it is not – but this is surely trumped by the way the illusion manifestly seems to us. Claims that all perceivers who are affected by such illusions are mistaken about their experience⁵⁹ must be adequately motivated, rather than just made, if they are to be convincing. It seems far from clear that DF's skill at coping in the world provides such motivation – we must keep in mind that she retains sensitivity to colour and texture cues which she can use to infer the sorts of visuomotor actions she might successfully prosecute, that she has had many years to develop strategies to deal with her condition, and that her exposure to hundreds of psychophysical tests has left her with a comprehensive sense of the capacities her spared perceptual sensitivity has left her with. We must also remember all that DF cannot do – draw or describe the form of an object, select an object from an array or use it based on its

⁵⁸ See Block (2007), and commentaries, for an argument that what we consciously experience far outstrips what we can report, as well as a swathe of empirical and conceptual arguments for and against this claim. Dretske (2006) and Clark (2007) also discuss these issues.

⁵⁹ Assuming that all such perceivers experience the illusions in the way I have described.

form, act on an object without being prompted (either by an experimenter or by her own inferential strategies based on colour and texture cues, or exploratory movements), match two visually presented objects based on their form, and so forth. I conjecture that if DF were (rather cruelly) required to perform novel but simple tasks in an unfamiliar monochrome and mono-textured environment, there would be far less temptation to suppose that she enjoyed unreportable visual experience of shape.

DF aside, the mere fact that processed information can be put to use in the guidance of certain kinds of behaviour does not strike me as adequate motivation for the claim that such information contributes to the character of a conscious experience. There are many sources of processing in the brain that have some kind of upshot in behaviour, and it is surely outlandish to suppose that each of these must make a contribution to conscious experience. Do I experience all the information which guides the countless subtle postural adjustments I make to keep my balance when riding a bike, or the information which leads me to shift from an uncomfortable posture when I am asleep? How about the semantic information by means of which I am (apparently) unconsciously primed to act in a variety of different ways in experiments such as those described in Chartrand and Bargh (1999)? Do blindsight subjects experience the information that guides the forced-choice discriminations they can make? It is possible, if perhaps counterintuitive, to answer 'yes' to each of these questions. But unless we wish to answer 'yes' to them all, there needs to be some rationale behind our contention that some bit of action-influencing information processing is reflected in our experience. My first response to the objection, then, is that the *prima facie* reasons to suppose that the action-guiding information processed by the dorsal stream is not part of our visual experience are good enough to place the burden of proof on the objector to come up with convincing motivations for the contrary claim⁶⁰.

Secondly, even if we were to suppose that the processing in the dorsal stream is reflected in conscious experience, it seems unlikely that the content of the dorsally-enabled portion of experience would support anything we could naturally describe as 'experience of visual form'. As Matthen (2005, p.301-305) argues, we need not suppose that the dorsal stream codes for the entire shape of an object. Whilst such shape information may be necessary for the selection of grip type and force, both features our discussion in section two gives us reason to think are handled by the ventral stream, to scale finger-separation in a manner appropriate to grasp some object requires only that the positions of the points to be grasped

⁶⁰ Actually, I think the action-space approach is in a stronger position than this. As we will see in subsequent chapters, there are independent reasons for adopting an action-space approach to conscious perception – in light of these reasons, the action-space account can provide a rationale for this interpretation of the DVS data.

on the object be coded for, relative to the positions of the fingers that will do the grasping. Fulfilling the functions of motion-guiding vision does not require that this information be integrated into a representation of the object's overall shape – indeed, to do so when such integration has already been achieved by the ventral functions that select which dorsally-governed action should be performed would be needlessly computationally profligate⁶¹. The conjecture that only specific aspects of the target object are coded for by the dorsal stream has empirical support from Ganel and Goodale (2003), who demonstrate that the width and length of a rectangular figure are processed separately when grasping it, but not when perceptually classifying it, and Schenk and Milner (2006), who demonstrate that DF is aided in shape-classification tasks when engaged in a concurrent reaching task only when the shapes to be discriminated were of different widths, suggesting that the dorsal information DF appears to be using is sensitive only to width, not shape⁶². Thus, were the information processing in the dorsal stream to issue in conscious experience, the content of that experience might be something like 'it is currently appropriate to grasp with a given finger-configuration in a given region of egocentric space'. But we have no reason to suppose that dorsal processing could disclose any more specific information about objects or their properties. It strikes me as quite difficult to imagine what such sparse phenomenology might be like – but in any case, it surely falls short of the unnoticed experience of visual form hypothesised by the objector⁶³. My second response to the objection, then, is that the sorts of experiences which might plausibly be enabled by dorsal activity are of marginal importance at best in understanding the content and character of our visual experience. It would be misleading to suggest that the dorsal

⁶¹ See Matthen (2005, p.303) for an argument to the same conclusion from an evolutionary perspective.

⁶² Both these findings look compatible with the Matthen's conjecture that only the relative positions of grasping fingers and points-to-be-grasped on the object are coded for by the dorsal stream, as well as the authors' hypothesis that only width is coded for. Schenk and Milner's result has the water-muddying implication that dorsal activity *can* effect classificatory abilities – I won't attempt to fully interpret this finding here, but note that DF still performs with far less accuracy (about 75%) than normally-sighted control subjects (100%), and her classifications are described as guesses, rather than reports.

⁶³ According to the action-space account we are developing, *if* the dorsally-processed width information were available to inform ongoing practical reasoning, then that information would figure in the content of a conscious experience. As I understand them, Schenk and Milner's (2006) experiments leave it underdetermined whether this is so. Whilst dorsally-processed width information can be made available to influence DF's classifications of shape, it is not clear that it can do this automatically and non-inferentially, in the way that the perceptual sensitivity of standard perceivers influences their classificatory behaviour – it seems that DF must be externally prompted to at least form an intention to initiate a reaching movement before her spared sensitivity can be put to any use in the classification task, suggesting that the dorsal information would not be conscious on the action-space account. I remain neutral here on just how these results are best interpreted, though.

stream could enable experiences just like those enabled by the ventral stream except in that they are unreportable. We might thus remain neutral on whether any dorsal information has an upshot in conscious experience, since the sorts of experiences the dorsal stream might plausibly enable were surely not a large part of the visual experiences that are our explanatory target. To clarify, consider Matthen's (2005, ch.13) plausible suggestion that the operations of motion-guiding vision are responsible for the 'feeling of presence' that visually-perceived objects have. The action-space account can agree with this if the dorsal information concerning presence is apt to inform the subject's ongoing practical reasoning. But even if this information turned out not to be apt in this way, and was thus a counterexample to the action-space view (and I see no reason to believe that it should – I raise the possibility only to illustrate the point), such dorsally-enabled content seems a marginal enough feature of visual experience that we could still claim that the action-space account was, by and large, correct.

Conclusion

We began this chapter with a brief review of sensorimotor theories of perception, then surveyed the evidence for the dual systems theory of visual processing. I argued that the DVS evidence gives us good reason to think that conscious vision is especially implicated in providing information about the enduring and objective properties of objects, and tied to capacities for memory and action-planning, each in a way that motion-guiding vision is not. Next, we considered the compatibility of DVS and the sensorimotor view, and I argued against several different attempts to demonstrate the direct incompatibility of the two theories – Jacob fails to demonstrate that sensorimotor theory is committed to an action-guiding role for conscious vision that would place it in tension with DVS (as does Block, though this point is less central to his treatment), Block's DVS-inspired arguments against sensorimotor theory are inconclusive, and Clark focuses his criticism on a feature of sensorimotor theory of which we had independent reason to be suspicious, and which we can consistently reject. However, in section four, I argued that DVS has stronger affinities with the action-space theorist's contention that perceptual experience consists in a grasp of the practical consequences of perception than the sensorimotor theorist's contention that perceptual experience consists in a grasp of the perceptual consequences of action. I argued that when viewed in the context of the debate between these two theories, we see that the real challenge to sensorimotor theory from DVS is the possibility that a grasp of sensorimotor relations is of only instrumental importance in an account of perception – important insofar as it contributes to a perceiver's grasp of the practical consequences of

perception. Finally, we considered the objection that our inferences from the DVS data to conclusions about the functions of conscious vision were too hastily drawn. I argued that we have *prima facie* reasons to trust our interpretation of the data that suffice to put the burden of proof on those who would argue that the processing of the dorsal stream can have an upshot in conscious experience, and that the sorts of experiences the information-processing of the dorsal stream could plausibly support form, at best, a marginal subset of what we standardly mean by 'visual experience'.

The conclusions of this chapter, then, are that the DVS data suggests that conscious vision involves an ability to integrate ones perceptual sensitivity to the objective properties in one's environment with reasoning and planning, and that this data gives us empirical reason to favour the action-space approach over the sensorimotor view. Next, I consider some puzzling cases in the domain of colour-perception, and argue that we can draw the same conclusions there.

Chapter 3: Special Spectacles, Synaesthetic Sensations and the Action-Space Theory

In the last chapter I argued that the DVS data meshes well with the action-space theorist's view of the relationship between action and conscious perception as sketched in 1.5 and 1.6. Thus, if we wish to use the DVS data to motivate some form of enactivism, we should favour the action-space view over the sensorimotor view. Moreover, if we can provide independent reasons for adopting an action-space theory of conscious vision then the empirical support that DVS affords the action-space view will be strengthened – the action-space account will be an independently plausible theory of visual experience that can predict and explain the dissociations uncovered by the DVS data, rather than one among several possible views of conscious vision that are consistent with that data. This chapter makes a start on providing such reasons⁶⁴. It is divided into two main sections. In the first, I argue that the action-space theory is better equipped to deal with certain puzzle cases than the sensorimotor view – firstly, adaptation to colour-distorting goggles, and secondly, the perception of synaesthetic sensations. I begin by going over the key features of Pettit's (2003) account of colour looks, drawing parallels with the action-space approach, then argue that such an account can deal with each case, unlike the sensorimotor view. As with our discussion of the DVS results, our overall moral will be that sensorimotor dynamics look to play at best an instrumental role in an account of perceptual experience. The second section pools the resources of the discussions of these cases and the results in the last chapter to formulate our first pass at an action-space theory, and responds to a range of obvious objections to such a proposal.

I.

3.1 Sifting, Sorting and Seeing

We briefly discussed Pettit's (2003) account of colour looks in 1.5 and 1.6. It is time to give it a fuller treatment. Pettit's discussion centres around two competing accounts of what it is

⁶⁴ Subsequent chapters continue this task – chapter four consists in a detailed discussion of what the action-space account should say about colour experience. Chapter five clarifies a key feature of our account – the requirement that the actions which a perceiver takes to be enabled must be apt for integration with her practical reasoning to issue in conscious experience. Finally, chapter six argues that the action-space account suggests how we might tackle the problem of the explanatory gap with which we began in chapter one.

for an object to look red (or any colour)⁶⁵. Firstly, according to what Pettit calls ‘the qualia theory’, an object looks red (has a red look) to a perceiver iff it produces a red quale in the perceiver – that is, iff it effects a state in the perceiver such that:

- i) The way the object looks to the perceiver is independent (at least in principle) of the perceptual abilities typically enabled by an object’s having such a look, such as sifting that object from differently-coloured objects, sorting it into the same category as similarly-coloured objects and tracking it over time, all on the basis of its apparent colour.
- ii) The way the object looks is ‘manifest’ for perceivers – they experience the object as looking that way.
- iii) The object manifestly enables the above perceptual abilities – perceivers see the look as “inviting relevant contrasts and comparisons: they make those contrasts and comparisons as responses that the look extracts from them; they don’t just find themselves disposed to make them, they know not why.” (Ibid, p.226)⁶⁶

Such an account has its merits – points ii) and iii) (or perhaps a slightly weakened version of the third point) appear to be truisms about colour perception, and qualia theorists will argue that i) looks like a truth about our experience – we might have the intuition that it is possible in principle to have an experience with the phenomenal quality of redness, but in the absence of the discriminatory abilities that are usually characteristic of such an experience⁶⁷. However, as Pettit notes, we might be suspicious of the fact that such a theory appears to rule out the compatibility of the ‘minimal naturalism’ with which we began chapter one and our account of colour experience. As Lewis (1995) argues, it seems plausible that the folk-psychological conception of qualia⁶⁸ entails that they are known in their essence by their subjects – thus, if a subject’s quale is subserved by a particular neural signature but the subject knows nothing about that signature, then the signature is, at best, an incidental property of that quale. Given that a subject can have a quale without

⁶⁵ Pettit is concerned with a sense of ‘looking red’ that reflects the way the perceiver takes the object to appear (not necessarily the way the perceiver takes the object to in fact be coloured) (see Ibid, p.221-222).

⁶⁶ It seems likely that a defender of qualia would question whether this is a necessary condition on a subject having a red quale (Pettit in fact says only that looks are ‘almost always’ credited with the three properties above by qualia theorists (p.226)). However, they should agree that this condition is typically met by perceivers.

⁶⁷ Or to have such abilities without the associated experience – recall the zombie and inverted spectrum arguments from 1.1.

⁶⁸ Insofar as there is such a thing – see Lewis (1995) for a very nice discussion of these issues, on which the present paragraph draws.

knowing anything about its physical or functional underpinnings, such underpinnings are all inessential properties of that quale. And, since the first point of our definition above entails that qualia are essentially independent of their typical functionally specifiable consequences, qualia have been severed of all essential ties to features in terms of which we might give a naturalistically acceptable account of them⁶⁹. Additionally, Pettit notes that such a conception of qualia will make it impossible to resist the modal arguments we considered in 1.1 – so long as qualia are conceived as essentially independent of their physical underpinnings or associated effects, there will be no obstacle to conceiving of physically identical situations with differing distributions of qualia.

In light of the metaphysical commitments of the qualia theory, Pettit defends an alternative, dispositional account of an object's looking red to a perceiver. On Pettit's view, the explanatory direction employed by the qualia theorist is reversed. According to the qualia theorist, the quale by means of which the subject sees that an object looks red explains the subject's ability to sift, sort and track the object on the basis of its look. Pettit argues instead that the ways that perceivers are enabled to sift, sort and track an object by their perceptual sensitivity to the microphysical properties that subserve its colour fully explain the way the object looks to them. To motivate this possibility, Pettit invites us to consider the way a ball moving quickly through the air looks:

“Is there an effect-independent way it looks such that we might imagine that look being absent – that ‘fastish’ looks, as we might say – when we remain visually inclined to judge that the ball is going fast, or being present when we become visually inclined to judge that the ball is moving slowly? Surely not. All that happens is that the ball has a look that is essentially tied to the effect of inducing in you a judgement ‘It’s going fast!’ and, no doubt more primitively, certain reaching and ducking responses.” (Ibid. p.229)

Pettit's thought is that we should think of an object's looking red in this way. Thus, when an object looks red,

⁶⁹ I'm puzzled as to how best to interpret Pettit's thinking here. He endorses Lewis's account of the connection between qualia and anti-physicalism, but fails to say why his definition of qualia (points i)-iii) above) entails the thought that qualia are known to subjects in their essence. The reading of 'manifest' required for the natural suggestion that this falls out of ii) seems to create a tension between i) and iii). The most plausible solution available to Pettit seems to me to be to drop iii) as a necessary requirement. We should certainly agree with Pettit's conclusion, though – that the traditional way of construing qualia renders them incompatible with physicalism.

“The way it looks is exhausted in the character the perceived object has as something that manifestly gives rise to certain effects. [...] The object looks red so far as it manifestly enables you to sift and sort and track it in the red-appropriate manner, and so to make corresponding judgements; it looks red so far as you see it as extracting those responses from you.” (Ibid. p.229-230)

The intuitive response is that such a proposal gets matters backwards; that the abilities Pettit alludes to are not the basis of our visual experience – rather, those abilities are based upon that experience. But Pettit’s account, and the action-space account we are developing, does not credit this intuition. As Pettit clarifies,

“The idea is that perceptual exposure to a red object tunes us in a way that enables us to sift, sort, track and so on, and that this tuning occurs ‘behind our back’; it occurs in virtue of the brain’s responses to the perceptual input at a subpersonal level. What we become aware of in the experience of something red presupposes that those subpersonally induced responses are in place; the red look of the object is the look it has so far as the perceived object is registered as enabling – as having already enabled – those responses...” (Ibid. p.232)

Pettit is not absolutely clear about what he takes to be the importance of such an appeal to the automatic, subpersonal enabling of the abilities he wishes to emphasise⁷⁰, but I take its utility to lie in undermining the supposed obviousness of the intuitive response above. Presumably, such a response is motivated by reflection on our experience – it seems introspectively obvious to us that our dispositions with respect to colour are grounded in our colour experience, not vice-versa. However, once it has been emphasised that the dispositions to be appealed to are automatically and subpersonally enabled ‘behind our back’ (as opposed to being the sort of considered and reflective judgements about how things look to us that we can and do make), it is no longer clear that the question of whether the enabled abilities ground our experience or vice-versa can be easily settled by introspection.

⁷⁰ He goes on to say that this is in accord with much of what we know about how the brain works – that perceived objects can have effects on us without our knowing it. However, he illustrates this with the DVS finding that perception of certain illusions enables accurate grip-scaling which conflicts with what subjects claim to know about the objects being grasped, concluding that the way perceiving an object enables us to act is not always routed via our experience. Whilst this is correct, his choice of example looks unfortunate in that the enabled grip-scaling abilities here have no apparent upshot in conscious experience, and so look like a counterexample to the approach Pettit advocates. One way in which the action-space account under construction here goes beyond Pettit’s approach is in suggesting an account of why some enabled actions have an upshot in experience whilst others do not.

Perhaps the intuitive position is that our experience enables us to sift, sort and track objects on the basis of our experience – but how would it seem, introspectively, if the appropriate enabling of those abilities were all there was to our having a certain experience, as Pettit suggests? If this were the case, we would have the same experience, that poised us to perform the same range of perceptual discriminations and behaviours, but without necessarily judging explicitly that those are discriminations and behaviours currently enabled for us. It doesn't seem obvious to me that there should be a clearly introspectible difference between these options. The intuitive response, then, does not give us sufficient reason to discount the possibility of Pettit's proposal.

So, if Pettit's proposal is to be rejected, it must be rejected for other reasons. Perhaps it is implausible to think that appeal to a range of enabled abilities could ground or explain our experience, either because the abilities appealed to are too sparse and simple an explanans to plausibly explain the explanandum of phenomenal experience, or because they are poor candidates to meet the constraints on a solution to the explanatory gap we set out in chapter one. Or perhaps there is insufficient reason other than the metaphysical extravagances of the qualia theory to adopt such a view. Both such points are addressed in the rest of this thesis – the question of how the action-space account (of which I take Pettit's position to be an instance) can provide a solution to the explanatory gap is addressed in detail in chapter six. In the next two sections, I try to give independent reasons for adopting Pettit's view of the relationship between action and perception by showing that it can give plausible accounts of two puzzling cases of colour perception.

3.2 Special Spectacles

The first case I want to consider is one that Pettit cites himself as support for his view. Interestingly, it is also cited by Hurley (1998) and Hurley and Noë (2007) in support of their sensorimotor positions. I will argue that the case supports the action-space view, but not the sensorimotor theory. Kohler (1964) conducted a series of experiments involving adaptation to colour-distorting goggles. In these experiments, subjects wore goggles with vertically-bisected lenses, each of which had a blue-tinted left half and a yellow-tinted right half. Upon initial donning of these goggles, subjects' colour experiences and their colour naming and categorising behaviours were predictably disrupted. A uniformly white wall would appear half blue and half yellow when the subject looked directly at it, or completely yellow or blue when looked at through the appropriate half of the goggles. However, after several weeks of wearing the goggles, subjects' experiences and colour

categorisations returned to normal – the distorting effects of the goggles had somehow been compensated for. For example, one subject reports that, by the 36th day of wearing the goggles “Even though a great variety of yellow and blue stimuli are transmitted by my spectacles and keep impinging on my fovea, I no longer experience the corresponding colour sensations,” and that by the 46th day, “If I look first at the blue part and then at the yellow part of my visual field, the latter does not increase in intensity no matter how long I have looked at the former” (Kohler (1964), p.111-112). The subjective reports of participants were borne out by a test requiring the subjects to adjust a colour wheel to a neutral shade of grey, using a button that varied the colour of the wheel between shades of blue, yellow, and neutral grey. The initial donning of the goggles predictably disrupted the subjects’ abilities to perform this task, but after adaptation they could perform the task with similar proficiency to their attempts before the goggles were put on (Ibid, p.106, 114). When the goggles were removed, subjects experienced the opposite sort of distortion in their colour perception to that which occurred when the goggles were first donned – in this case, the left half of the visual field was experienced as tinted yellow, and the right half as tinted blue. For example, Kohler reported that “Whenever I open up a book, I am amazed to find that the left page looks yellow, and that it looks white with a bluish tinge as soon as I turn it over to the right,” and that “When I take a walk, I am always conscious of a peculiar glare on my left, as if someone carrying a lighted candle were accompanying me; on my right, nothing of this kind occurs” (Ibid, p.115). Over time, this distortion also fades away, and subjects’ colour perceptions return to normal.

What do these results tell us about the relationship between action and perception? Hurley and Noë (2007) attempt to account for these results on behalf of the sensorimotor theory by claiming that:

‘The sensorimotor expectancies characteristic of particular colours relate ultimately to the underlying invariant patterns of dependency of sensation on movement, and these do not change when the goggles are worn. But they are given new clothing, a transformed implementation, and as a result the perceiver’s understanding of them is disrupted until his expectations have adjusted to this new implementation and related it to the underlying invariant patterns.’ (Hurley and Noë (2007), p9-10)

The idea is that the sameness of experience before the goggles are donned and after adaptation has occurred is explained by the subject’s sensitivity to an underlying invariant sensorimotor dependence between perception and movement. The goggles disrupt this

sensitivity by giving this dependence a ‘new implementation’ which is filtered out, or compensated for, over the course of the adaptation. So, to adequately account for these results, the sensorimotor theorist must make it plausible that colour perception is a matter of understanding some kind of dependence of perception upon movement, and that, at some appropriate level of description, that dependence is largely identical before the goggles are donned, and after the goggles have been adapted to, thus explaining the similarity between pre-goggles and post-adaptation stages of the experiment. They will explain the distortion in colour experience by pointing towards the obvious distortion in the dependencies between perception and movement once the goggles are donned (saccading right makes everything appear tinted yellow, saccading left makes everything appear tinted blue, and keeping eyes still whilst moving one’s head makes a vertical partition between yellow and blue-ish portions of the visual field sweep along with one’s head movements), but claim that these altered sensorimotor dependencies are somehow ‘superficial’ – perceivers learn to see beyond them, to the invariant sensorimotor dependencies that obtain through all pre- and post-goggle stages of the experiment, and their colour perception adapts to the extent that they manage to do this.

I suggest that sensorimotor theorists face a problem when attempting to specify exactly what the relevant sensorimotor invariant is. To see this, first recall (from 2.1) that there is an ambiguity in the appeal to the dependency of perception on action - the ‘perception’ in this dependency admits of a personal-level and a subpersonal-level construal⁷¹. For example, a sensorimotor theorist might construe the perceptions that vary according to our movement either as subpersonal activity (such as patterns of retinal stimulation, or at some higher level of visual processing) or personal level visual experience. But neither such construal can provide the invariant sensorimotor dependence required for an explanation of Kohler’s results.

Sensorimotor relations which obtain between perceptual experience and movement are not invariant, since these relations clearly change when the goggles are donned (everything looks blue when the subject looks left, yellow when she looks right), return to normal over the course of adaptation, and change again when the subject removes the goggles. If we restrict our attention to the sensorimotor relations pre-goggles and post-adaptation, we find an invariance, but this ignores the disruption and recovery of the original experiential sensorimotor relations which occurs during adaptation, and the second disruption which occurs once the goggles are removed. Specifying the way in which these relations are disrupted, recovered, and disrupted again might be a partial *description* of the subject’s

⁷¹ As, we noted, does the ‘action’, though this won’t concern us here.

experience through the different stages of the experiment, but it does not give us an example of a sensorimotor invariant which we need to appeal to in order to *explain* the adaptation. So, a sensorimotor invariant cannot be found in the relations between perceptual experience and movement, since the relationships between the way things look and the perceiver's movements differ over the different stages of the experiment.

Neither, on the most obvious construal, are sensorimotor relations which obtain between subpersonal stimulation and movement invariant, since donning the goggles introduces a new dependency between eye and head movements and systematic shifts in the wavelength of light hitting the retina. This new dependency continues to obtain after adaptation has occurred, but the subject's experience has reverted to the way it was when the normal set of dependencies was in place. On this construal, we have an invariance in the stages of adaptation after the goggles are donned. This might seem heartening for the sensorimotor view – it could be that the adaptation process constitutes gradually latching on to this new invariance. But remember that the challenge was to come up with a pattern of dependency that was invariant across pre-goggle, adaptation, and post-adaptation periods. Our second construal fails to provide an invariance between the pre and post-goggle-donning stages. A sensorimotor invariant cannot be found in the relations between retinal stimulation and movement, since these relations change when the goggles are donned, remain the same over the course of adaptation, and change once again when the goggles are removed.

Perhaps, then, the sensorimotor theory should focus on relations between movement and subpersonal goings-on at a level higher than that of retinal stimulation, such as patterns of cortical stimulation. However, the subpersonal processes underpinning colour perception are multifarious, and can be described at many different levels – thus, claiming that Kohler's results should be explained by reference to the relations between those processes and movement leaves the content of this solution radically underdetermined.

Additionally, it would surely be surprising if we found that the relations between some portions of subpersonal processing and movement were straightforwardly invariant throughout the stages of the experiment – visual colour processing consists in transforming information in the light hitting the retina into information about the colours in the world around the perceiver, and when systematic distortions in the relations between the character of the light hitting the retina and movement are introduced, we would expect the character of any processing that codes for those relations to be similarly distorted. The suggestion that some such processing might remain invariant over the different stages of the experiment would seem to entail that there was some pattern of activity that helps

subserve colour perception that remains unchanged even when the light hitting the retina is dramatically distorted. The possible function of such processing would be unclear – why would the brain employ a coding scheme for colour that was neutral as to whether the subject was or was not wearing (a certain type of) colour-distorting goggles?

It seems more likely, then, that any invariant relations between subpersonal processing and movement would only be found by focussing on a higher level of abstraction than cortical activity or the mechanics of subpersonal processing. Perhaps the sensorimotor theorist should hold that the relations between subpersonal processing and movement *are* systematically distorted when the goggles are donned, but that there is some higher-order similarity between the original set of relations and the relations reflecting the distortion introduced by the goggles. If this is the possibility that the sensorimotor theorist wishes to pursue, then we should note firstly that more needs to be said about what this higher-order similarity might be if the sensorimotor theorist claims to have accommodated these results. For there are indefinitely many ways in which the processing/movement relations that obtain over the different stages of the experiment might be said to be similar. They are similar in that they are both sets relations between events in the brain and movements of the body, and in that they are both mentioned in this section of this thesis, but clearly higher-order similarities at these levels cannot ground an explanation of Kohler's results. The sensorimotor theorist thus needs to say more about what kind of higher-order similarities they have in mind, and why they are more promising candidates than those just mentioned. Secondly, we must bear in mind that if the explanation of Kohler's results is to count in favour of sensorimotor theory, the higher-order similarity must be given in sensorimotor terms. That is, the property that unites the relations between processing and movement over the separate stages of the experiment must be the property of perception varying with movement in a certain way. If the common property were instead (for example) that each set of sensorimotor dependencies allows the perceiver to represent the colours in their environment as being distributed in a certain way, then our theory of colour perception would turn out representationalist rather than sensorimotor – the importance of sensorimotor relations would be merely instrumental, consisting in enabling a certain representational state in terms of which the content and character of the perceiver's experience was then explained. It seems, then, that the only way the sensorimotor theorist might account for Kohler's results is to specify a set of high-level dependencies between perception and movement that is shared by the various sets of lower-level dependencies between perception and movement over the different stages of the experiment. However, it is opaque to me what sort of high-level dependencies these might be – I have just argued that they cannot be dependencies obtaining between movement and experience, retinal stimulation or cortical activity, and nowhere in the

development of the sensorimotor theory (to my knowledge) do other possible candidates suggest themselves. It seems, then, that the invariant feature needed to make sense of Kohler's results cannot be specified in sensorimotor terms.

Whilst it proves difficult to even make a case for the possibility of an account of Kohler's results in terms of dependences between perception and movement, an account of those results in terms of the abilities enabled for perceivers by their colour perception suggests itself fairly naturally, as Pettit notes⁷². According to an action-space theory, what is disrupted and restored over the course of the goggle experiment is the range of actions and abilities enabled for the perceiver by coloured objects. Before the goggles are donned and after adaptation occurs, an identical space of colour discriminations, categorisations and judgements are enabled by the subject's perceptual exposure to a coloured object. The invariance that allows the subject to adapt to the disruption of these abilities caused by the goggles is, on this account, the fact that coloured objects are objectively apt to be sifted, sorted, tracked and otherwise categorised and acted upon on the basis of their colour by the perceiver in the same way throughout the stages of the experiment. The goggles disrupt the perceiver's sensitivity to this fact by introducing a new set of sensorimotor dynamics. Adaptation consists in compensating for these altered dynamics to bring the perceiver's range of colour-related dispositions and intentions – the space of actions elicited from the subject by exposure to a coloured object – back into line with the way in which colour properties are actually distributed in the world. The degree of similarity between the colour experiences of a subject wearing the goggles and her experience before donning the goggles is determined by the degree of similarity between the sets of abilities enabled for her before and after she donned the goggles. As we noted, subjects who have adapted to the goggles can match and categorise colours with similar proficiency to before the goggles were donned – according to the action-space account, the restoration of the set of abilities characteristic of colour perception is all there is to the restoration of veridical colour experience. The 'intuitive response' to this proposal, considered in the last section is that it gets matters backwards – these abilities are restored because of the adaptation of subject's experience, not vice-versa. But, as argued previously, our introspective evidence is neutral between the 'intuitive' and action-space approach, so long as we bear in mind that the abilities we are appealing to are enabled automatically, 'behind the subject's back', as opposed to being considered judgements by the subject on their experience. Considering Kohler's results seems to provide support the action-space view over the intuitive view

⁷² Ibid, p.244-246. Pettit does not go into much detail about what the invariant feature needed to explain the results could be – what follows is my own view on behalf of the action-space theorist.

that colour categorisations are made on the basis of essentially effect-independent qualia – if we suppose instead that colour experiences are essentially tied to the abilities associated with them, the concurrent adaptation of colour experience and colour categorisation behaviour can be readily explained. Conversely, as Pettit notes,

“Short of positing an independently unexplained mobility in *qualia*, the *qualia* theory would naturally predict that after wearing the glasses for a suitable period, people would report that whilst the bluish and yellowish tints introduced by the glasses remained in place – whilst the distorted *qualia* remained in position – they had now learned to be very good at allowing for the change in making judgements on the colour of things; they had learned to recondition their judgemental responses to the altered *qualia*.” (Ibid. p.244-245)

Moreover, our discussion suggests that another moral to be drawn from Kohler’s results is that knowledge of sensorimotor relations is of only instrumental importance in explaining the content and character of visual experience. Implicit knowledge of sensorimotor relations might be part of what is involved when we come to know the nature and extent of the actions our perceptual sensitivity to our environment puts us in a position to perform. But Kohler’s results imply that a range of very different sensorimotor backdrops are consistent with a perceiver’s ability to grasp that she is in such a position. Those results thus suggest that the character of experience is fixed by the space of actions that are perceptually enabled for a perceiver at a time, not by our familiarity with whatever sensorimotor dependencies such enabling may involve. The way things look to perceivers over the stages of Kohler’s experiment reflects what they take themselves to be able to do on the basis of their perception, not what they know they could perceive as a result of their actions⁷³. We should conclude, then, that Kohler’s results support the action-space approach to perceptual experience.

3.3 Synaesthetic Sensations

⁷³ Note that this section has not aimed at a full account of Kohler’s results on behalf of the action-space approach. Doing so would require an of the invariant feature which explains the adaptation, and of the abilities which the experiences enabled by perceptual sensitivity to that feature are claimed to consist in. The further discussion of colour experience in the next chapter goes some way towards fulfilling these aims. However, for my purposes here (adjudicating between the action-space and sensorimotor views), the above outline of an action-space treatment of Kohler’s results will suffice.

The next puzzle case I want to consider is synaesthetic colour experience. Like Kohler's results, this case is discussed by Hurley and Noë (2007) – however, they argue there that synaesthetic experience presents a potentially fatal counterexample to both sensorimotor and action-space approaches to perceptual experience. I want to suggest that eschewing an emphasis on sensorimotor relations in favour of an emphasis on perceptually enabled abilities suggests how we might give an account of synaesthetic colour perception, and thus that synaesthetic colour experience constitutes another case that suggests we should prefer the action space approach to the sensorimotor view.

Synaesthesia is a condition where veridical sensations of one type are accompanied by sensations of another type – for example certain sounds, words or names can be accompanied by sensations of taste⁷⁴. In colour-grapheme synaesthesia, words or numbers are associated with sensations of specific colours. In colour-phoneme, or coloured-hearing synaesthesia, certain sounds are accompanied by sensations of colour (photisms). The manner in which photisms are experienced varies with different subjects and different varieties of synaesthesia. Some synaesthetes report experiencing a generic coloured shape, or a coloured visual image of the word or letter being seen or heard. The photism may be experienced as non-localised, overlaid upon or filling a visually presented letter, “bound to the sound” of an aurally presented one, or located upon an invisible plane located at arms reach (Grossenbacher and Lovelace (2001)). Three points about synaesthetic experiences are particularly important for our purposes here. Firstly, synaesthetic sensations are genuine features of the experience of synaesthetes (as the results to be considered below illustrate) – genuine synaesthetes are not speaking metaphorically, or delusional, when they claim that they experience (for example) letters as accompanied by specific colours. Secondly, synaesthetic colours are experienced as of the same phenomenal type as veridical colour sensations – synaesthetic green is experienced as phenomenally similar to veridical green, rather than somehow *sui generis*. Thirdly, synaesthetes nonetheless automatically know that their synaesthetic experiences aren't real – they do not mistake synaesthetic experiences for experiences of real, worldly coloured things.

The problem posed for the sensorimotor theorist by synaesthetic colour perception is that whilst, for the sensorimotor theorist, phenomenal looks are ultimately explained by reference to our sensorimotor knowledge and expectancies, synaesthetic colour looks seem to

⁷⁴ The opening lines to Bob Dylan's 'Chimes of Freedom' are evocative of the running-together of different modalities that characterises synaesthesia: “Far between sundown's finish, and midnight's broken toll/We ducked inside a doorway, thunder crashing/As majestic bells of bolts struck shadows in the sounds/Seeming to be the chimes of freedom flashing”.

constitute an example of similar phenomenology⁷⁵ (that of, e.g., veridical green and synaesthetic green) with different sensorimotor knowledge and expectancies implicated in each case. The sensorimotor theorist wishes to explain the phenomenology of a veridical perception of a green object in terms of our sensorimotor knowledge of how our sensations will change as we move around it, expectations about how sensations would change in a variety of similar and different lighting conditions, and generally, the range of sensorimotor skills implicated by our perception of the coloured object. This strategy does not seem to work for (e.g. coloured-hearing) synaesthetic sensations⁷⁶. Firstly, talk of sensorimotor skills implicated by a synaesthetic sensation seems misplaced, since there is no worldly object for the subject to perceptually engage and interact with. More importantly, synaesthetes lack the kind of expectancies about how their photism will change as they move around and as lighting and other perceptually relevant conditions differ in terms of which sensorimotor theorists might naturally explain colour experience⁷⁷. Rather, synaesthetes are aware that the expectancies appropriate to worldly colours do not apply to their photisms, and the fact that they never mistake photisms for worldly colours suggests that this awareness is automatic and non-inferential. It would thus seem implausible for the sensorimotor theorist to respond that synaesthetes *implicitly* take the same range of expectancies to be appropriate to synaesthetic and veridical colours, whilst realising on an intellectual level that this is not so – were this the case, we would expect synaesthetes to occasionally absent-mindedly mistake photisms for worldly colours, which they do not appear to do. It seems plausible that even the synaesthetes' implicit sensorimotor expectancies concerning the behaviour of synaesthetic colours differ from the corresponding expectancies concerning worldly colours.

So, because of the lack of commonality between the sensorimotor profiles of the two phenomenologically similar types of sensation, it seems that the sensorimotor view lacks the resources to explain this similarity. Thus, Hurley and Noë concede:

⁷⁵ It seems that there must be *some* difference in phenomenology, since synaesthetes automatically differentiate synaesthetic colour sensations from veridical ones. However, as we see below, the challenge for the sensorimotor view is to account for the *commonalities* between synaesthetic and veridical experience, rather than the differences.

⁷⁶ Hurley and Noë put the challenge to their account in terms of coloured-hearing synaesthesia, but the problem generalises to any synaesthetic experience where the subject apparently fails to have sensorimotor expectations concerning their synaesthetic sensation. In the positive story I give on behalf of the action-space view, I focus chiefly on colour-grapheme synaesthesia, since that is the form most widely discussed in the empirical literature. In this and the next chapter I focus only on the way the action-space account should treat colour experience – however, in the concluding chapter I briefly suggest how the action-space account could be extended to

⁷⁷ The next chapter contains a more detailed discussion of possible sensorimotor approaches to colour perception.

“...in the case of synaesthesia, pretty well the whole set of hunter-gatherer powers and sensorimotor expectancies associated with normal colour looks is missing. Not only can synaesthetic colours not actually be used to guide hunting and gathering, but synaesthetic colour experience does not generate expectations that sensations will vary with movement, with light, with background, or any of the normal sensorimotor expectations that perceptual experience of colour generates. It is a truly hard case for the hunter-gatherer perspective.” (Hurley and Noë (2006), p.22)

They conclude that further work, looking into the neural basis of synaesthetic sensations is required to determine whether the hunter-gatherer perspective can meet the challenge from synaesthesia.

However, whilst synaesthetic sensations constitute a damaging counterexample to sensorimotor views, I argue now that an action-space view is well-equipped to give an account of them. For although the sensorimotor relations and expectations characteristic of synaesthetic sensations differ markedly from those characteristic of veridical sensations, there seem to be significant commonalities in the range of abilities enabled by veridical and synaesthetic colour sensations. The empirical literature on synaesthesia makes a compelling case that synaesthetes’ abilities to sift, sort, track, categorise, re-identify, and otherwise use their photisms based upon their apparent colours are highly similar to the corresponding abilities enabled for perceivers by perceptual sensitivity to veridical colours. I want to suggest that an enactive theory can appeal to these commonalities to account for the similarity of synaesthetic and veridical colours, and thus that reflection on the case of synaesthesia gives us another reason to favour the action-space over the sensorimotor view.

The first and most simple example of such a commonality is the way in which **synaesthetes categorise and re-identify synaesthetic colours in the same ways as veridical ones**. Baron-Cohen found that the detailed verbal descriptions of synaesthetes’ photisms were 100% consistent over separate trials, compared to the 17% consistency of control subjects (cited in Smilek and Dixon (2002), p.2). Similarly, synaesthetes have shown consistency in matching the colour of their photisms to very specifically coloured chips over successive trials (Ramachandran and Hubbard (2003)). This suggests that synaesthetic colours stand in relations of similarity and difference to each other comparable to the relations between worldly colours – synaesthetes verbally describe and categorise the colours of their

photisms in the same way as they would worldly colours, and can place synaesthetic colours in relations of similarity and difference to worldly colours⁷⁸. As we shall see in the next chapter, such enabled abilities to place colours within a network of similarity and difference relations on the basis of perceptual sensitivity are a key feature of the action-space account of colour experience.

That capacities to classify and discriminate between synaesthetic colours are similar or identical to those implicated in veridical colour perception are involved in synaesthetic perception is also suggested by Ramachandran and Hubbard's (2001) demonstration of **pop-out effects for synaesthetic colours**⁷⁹. In these experiments, colour-grapheme synaesthetes are shown an array of "5"s and "2"s, with the "5"s forming a shape (such as a triangle) against the background of "2"s. Because the synaesthetes see the numbers as different colours, the triangle will quickly "pop-out" for them, whereas non-synaesthetes take significantly longer to discern the location and shape of the "2"s within the "5"s. This suggests that capacities to detect boundaries and contrasts in colour, and to appropriately integrate and utilise this information in perceptual activities, are present in synaesthetic experience as in veridical experience. Similarly, Palmeri et al (2002) found that the synaesthete they studied was aided in a visual search task when the target letter had a differently coloured photism to the surrounding objects⁸⁰.

Conversely, Smilek and Dixon (2002, p.9) have demonstrated **interference effects** for synaesthetic photisms. They showed that colour-grapheme synaesthetes are slower at identifying a black letter placed on a background of the same colour as its photism than when the background is of a different colour to the photism. This suggests that the abilities to easily sift a coloured object from an incongruently coloured backdrop but not

⁷⁸ Ramachandran and Hubbard (2001, p.5) claim that "Synaesthetes often report 'odd' or weird colours they cannot see in the real world but see only in association with numbers". The other work I cite here makes no mention of this fact, but the account of 'impossible' colours I give in the next chapter suggests how an action-space approach might account for this – I return to this point there.

⁷⁹ However, Edquist et al (2006) failed to obtain similar results – see Edquist et al for discussion. Thanks to Susan Hurley for bringing this work to my attention. A salient difference between the search tasks studied by Ramachandran and Edquist might be that Ramachandran demonstrated pop-out for a *pattern* of photisms against a background of distracters, whereas Edquist et al failed to find pop-out effects for a *single* photism against such a background. I won't attempt to settle this dispute here, though - if synaesthetic pop-out does not occur, the other results surveyed in this section seem to me sufficient to make a case for an action-space approach to synaesthetic photisms nonetheless.

⁸⁰ Though cf. Sagiv et al (2006), who report a study that suggests that synaesthetic photisms have no significant influence on performance in certain visual search tasks.

from a similarly coloured one is the same (or similar) for synaesthetic and veridical colour sensations.

Other similarities in SST-abilities are suggested by **stroop effects**. In the colour-grapheme synaesthetic stroop effect, when a synaesthete attempts to state the (veridical) colour of a target letter their performance is impaired if the target colour is incongruent with the colour of their photism (Rich and Mattingley (2002), p.47-8). A variant of this for coloured-hearing synaesthesia is Gray's "Alien Colour Effect" – the ability of coloured-hearing synaesthetes to name ink colours is disrupted when the sound of the name of that colour has a photism with a different colour (cited in Hurley and Noë (2007), p.12). The extent of these effects is what we would expect if a subject were to have a veridical sensation of a colour whilst trying to name an incongruent colour. Thus it seems that synaesthetic and veridical sensations can disrupt some of our perceptual abilities in the same way. Hurley and Noë reject the suggestion that this result could be used to defend the hunter-gatherer perspective, since "...it is not plausible to claim that naming powers and expectations alone unite normal and synaesthetic colour looks" (Ibid, p.23). We can agree with this, yet still appeal to the alien colour effect results as providing evidence for the similarity of enabled abilities characteristic of synaesthetic and veridical sensations. Firstly, *contra* Hurley and Noë, we are not appealing to disruption of naming effects as the *sole* common ground between synaesthetic and veridical sensations. Secondly, unlike the sensorimotor theorist, we are not committed to specifying the common ground between synaesthetic and veridical sensations in sensorimotor terms. The disruption of perceptual categorisation abilities which occurs during the alien colour effect is just the sort of commonality to which the action-space account can appeal⁸¹.

In the same way, the action-space account can appeal to **priming effects** – synaesthetes can state the colour of a patch presented to them more quickly if it is preceded by presentation of a stimulus with a congruently coloured photism. Incongruently coloured photisms impair their ability to give the colour of the patch (Smilek and Dixon (2002), p.7-8). Again, the action-space account suggests that this is not merely a disruption in verbal abilities, but of the perceptually-enabled abilities of synaesthetes to identify and classify stimuli. And again, the nature and extent of the disruption is the same as we would expect from veridical colour priming effects, suggesting commonalities between the sets of perceptually enabled actions characteristic of synaesthetic and veridical colours.

⁸¹ Note that we should not think of the alien colour effect as a mere disruption of *verbal* abilities. Rather it is a disruption of our *perceptual* abilities to discriminate the colour of an object.

In light of all these results, an empirical prediction of the action-space approach might be that the extent to which the colours of photisms enhance or interfere with synaesthetes performances in search or classification tasks, or to which priming or stroop effects are found is sensitive to the relations of similarity and difference between the colours involved in a comparable way to the corresponding effects with veridical colours. For example, pop-out effects will be less marked and interference effects more marked if there is a red photism against a background of pink distracters than if there is a red photism against a background of green distracters. Likewise, priming effects will be more pronounced the closer the colour of the priming photism is to the colour of the patch to be subsequently named. The action-space account explains the similarity of synaesthetic and veridical colour experience by appeal to the similarities between the sets of abilities characteristically enabled by each type of experience. Thus, the action-space account predicts that, to the extent that synaesthetic and veridical colours are experienced as similar, the differences in influence of different combinations of colours on perceptual abilities illustrated in the search, classification, priming and stroop results surveyed above will be similar for both veridical and synaesthetic colours.

Do some of these results support the action-space view more than others? It is perhaps worth noting that there is some controversy over whether the processing responsible for synaesthetic photisms is pre-attentive or not – does attention have to be directed upon a grapheme before its photism is experienced? Of the results surveyed above, Ramachandran and Hubbard's demonstration of pop-out effects, Palmeri et al's demonstration that incongruent photisms aid visual search, and some of the priming and interference effects demonstrated by Smilek and Dixon look to support the hypothesis that synaesthetic colours can enable certain abilities for the perceiver pre-attentively, whereas the other results are consistent with the hypothesis that attention is required before any perceptual abilities are enabled by synaesthetic colours. The latter hypothesis seems to be supported by Edquist et al (2006), who also survey other work that suggests this conclusion. The results that suggest that synaesthetic colours are registered pre-attentively might seem to provide stronger support for the action-space view – recall the intuitive response to the action-space view, that the abilities associated with a colour experience are enabled on the basis of that experience, rather than serving as the basis of that experience. If some of the abilities in terms of which the action-space account wishes to explain synaesthetic colour experience are enabled for the subject before they attend to the grapheme that induces the photism, it becomes more difficult to press this objection – if the abilities in terms of which synaesthetic colours are to be explained are enabled

automatically, independently of the subject's awareness of the features that enable them, then it cannot be objected that those abilities are enabled quasi-inferentially, on the basis of the subject's awareness of such features⁸². The findings that suggest that synaesthetic colours are registered pre-attentively therefore look particularly good for the action-space view.

However, an action-space account of synaesthetic sensations would still be tenable were it to turn out that attention is required for the perception of synaesthetic colours⁸³. The important point for the action-space view is that the abilities in terms of which we wish to explain the commonalities between synaesthetic and veridical colour experience are not enabled on the basis of an already-existent colour experience. As just noted, this seems especially plausible if those abilities are enabled before the subject becomes aware of the relevant stimulus. But even if the stimulus must be attended before a synaesthetic photism is experienced, it does not follow that the enabled abilities characteristic of that experience are dependent on it, rather than vice-versa. Suppose that the relevant characters in the display used by Ramachandran and Hubbard to demonstrate synaesthetic pop-out must be attended before pop-out occurs – this is still consistent with our action-space approach, which could then claim that attention is required before the enhanced sifting and sorting abilities characteristic of the colour pop-out effect are enabled for synaesthetic subjects. The action-space account would be undermined if it were compelling to suppose that all the abilities evinced by synaesthetes in the experiments discussed above were enabled only by an existing synaesthetic experience – but it the results we have considered in this section do not seem to warrant this supposition⁸⁴.

I conclude that there is evidence suggesting substantial commonalities between the sets of actions enabled by perception of synaesthetic and veridical colours which can be used to

⁸² This point need not rely on the contested (though in my view defensible) assumption that we only experience what we attend. It is fairly uncontroversial that attention is required for a perceived feature to be available as the subject of an inference, or control of intentional behaviour. Thus, the claim that photisms enable their associated abilities only *inferentially*, on the basis of their experienced colours does not square with the claim that those photisms (and the abilities in terms of which the action-space view wishes to account for them) are registered pre-attentively.

⁸³ Currently there seems to be no consensus on this issue in the literature – see Edquist et al (2006) for review and discussion.

⁸⁴ It might seem natural to suppose that the abilities to categorise and re-identify synaesthetic colours are enabled on the basis of synaesthetic experience, rather than vice-versa. As suggested by the previous sections of this chapter, the action-space account denies this, claiming instead that those abilities are a key part of the basis of the colour experience. I made the case that this is a tenable position in discussing the 'intuitive objection' to the view in 3.1, and in the next chapter we will see further arguments for adopting this view of colour experience.

explain their phenomenal similarity. This contrasts with the dissimilarities of the sensorimotor profiles of synaesthetic and veridical colours, and the corresponding inability of the sensorimotor theory to explain their similarity⁸⁵. Considering synaesthetic sensations constitutes further evidence that visual experience should be explained in terms of the abilities enabled for a perceiver by their perceptual sensitivity, not in terms of that perceiver's grasp of sensorimotor dynamics.

II.

Our discussion of the explanatory gap in the first chapter, the DVS results in the second, and the ways enactive accounts might accommodate the cases we focused on above give us sufficient materials to state our first pass at an action-space account of perception. Here I outline the view, and recap the evidence in favour of it that we have surveyed thus far. I'll also try to further illustrate and motivate the view by briefly responding to a range of obvious objections. These are grouped into two categories – objections that claim that the conditions set forth for the action-space account are not necessary for perceptual experience, and objections claiming that those conditions are not sufficient for perceptual experience. Bear in mind, however, that the case for the action-space view will continue to emerge over subsequent chapters, as the key features of the view and its contribution to a dissolution of the explanatory gap are clarified.

3.4 The Action-Space Approach

In 1.5 we defined an action-space account of perception as one according to which conscious perception is a matter of knowing how what one sees affects what one can do – of understanding the practical consequences of movement. We also noted, at that stage without argument, that a perceiver's sensitivity to the actions currently enabled for her by the environment must show up not merely in her behaviour, but also in her practical reasoning and action planning if that sensitivity is to issue in conscious experience. In 2.2, we saw that a natural interpretation of the DVS results supports this requirement – our discussion of the DVS data suggested that capacities to integrate perceptual information

⁸⁵ For what it's worth, perhaps the above discussion also suggests that our intuitions lie more with the SST than sensorimotor perspective. Despite the lack of similarity between sensorimotor profiles of veridical and synaesthetic sensations, we are still willing to take synaesthetes' reports at face value. However, I think we would be more skeptical if synaesthetes claimed to have rich synaesthetic experiences without any attendant skills characteristic of colour-perception (i.e. if priming and pop-out effects were absent, synaesthetes could not use their photisms to aid memory and cognitive tasks, etc.).

with practical reasoning and action-planning were distinctive features of conscious visual perception. The other moral we drew from the DVS data was that conscious experience seems especially involved in representing objective and enduring properties of objects, allowing for objects and properties to be identified, classified and compared in ways that abstract from certain details of the perceiver's current egocentric relations to them. In light of our consideration of Pettit's position in 3.1, and the ways in which it can accommodate the puzzle cases of 3.2 and 3.3, the action-space account makes the strong claim that the contents of conscious visual experience be understood in terms of the abilities enabled for an agent by their perceptual sensitivity to their environment. Just as I have been urging in this chapter that we need not think of the abilities characteristic of colour perception as based upon an essentially effect-independent colour experience, we should not think of the abilities to identify, co-classify, compare and so forth that DVS suggests are distinctive of experience as enabled on the basis of an essentially effect-independent perceptual state. According to the action-space view, then, conscious experience consists in a abilities to classify, compare, sift, sort and track certain objects and features in her environment being primed for a subject by her perceptual sensitivity to that environment and put in touch with her capacities for practical reasoning and action-planning. We can sum this up by saying that conscious experience is understanding the ways in which one is currently poised to act in one's environment, or *knowing poise over a space of actions*⁸⁶.

At this stage, the action-space account will doubtless strike the reader as in need of both further motivation and elaboration. Subsequent chapters take up both these tasks, but let us review the case that can be made for the action-space approach based on our discussion thus far. The minimal conclusion of our consideration of the DVS results in chapter 2, and of the puzzle cases in this section was that we should prefer an action-space theory to a sensorimotor account of perception when adjudicating between enactive views. But our discussion has also provided support for the action-space account in its own right.

We can see the action-space account as developed thus far as consisting of two claims: firstly, that the content of perceptual experience should be understood in terms of abilities enabled for the subject by their perceptual sensitivity to their environment; secondly that those abilities form the basis of a perceptual experience when suitably integrated with the subject's capacities for practical reasoning. The chief support for the first claim has been our discussion of Pettit's theory in this chapter, and the account it suggests of the cases considered in 3.2 and 3.3. The most obvious objection to such a claim stems from the intuition that enabled abilities are simply not the right kind of explanans

⁸⁶ I owe this locution to (some combination of the influences of) Andy Clark and Tom Roberts.

for conscious experience. This intuition might have a number of sources, about all of which the action-space account has something to say. Perhaps the most obvious is the claim that the abilities associated with an experience are enabled on the basis of that experience, rather than the enabling of those abilities serving as the basis of the experience. But as we noted at several points above, it is not clear that introspection should be able to cleanly adjudicate between these options. If this is so, the arguments for the action-space construal of experience must be assessed on their merits, not ruled out in advance. Another motivation for the objection might be the claim that the abilities enabled for an experiencing subject are simply too sparse and coarse-grained to explain something as rich and complex as visual experience. We shall see in detail that this is not so in the next chapter's discussion of colour perception, but we can make a preliminary response on behalf of the action-space account here: since the enabled actions we are appealing to include capacities to discriminate, sort and track, the grain of the abilities that might be enabled for a subject will coincide with the grain of the perceptual qualities to which the subject is sensitive. It seems natural to suppose that it only makes sense to ascribe perceptual sensitivity to some quality to a subject if the presence of that quality has some upshot in the possible discriminations and classifications enabled for that subject – it would seem unwarranted, for example, to say that a perceiver is sensitive to the difference between roundness and squareness if they showed no sensitivity to that difference in their thoughts and deeds concerning round and square objects, instead treating those classes of objects identically. For every difference in what a subject perceives, however slight, the action-space account holds that there is a corresponding difference, however slight, in the range of abilities currently enabled for that subject⁸⁷.

The second key feature of the action-space account – that the abilities enabled for a subject by their perceptual sensitivity must be integrated with practical reasoning if they are to issue in conscious experience – also bears on the first. It might be objected that the discriminations and classifications to which we the action-space account appeals are not 'actions' in a sufficiently robust sense of the term. Perhaps they should instead be construed as the automatic results of subpersonal processes – as informational states registering the presence or absence of a feature, rather than sets of enabled abilities. I would like to resist this suggestion, and hold that all perceptual states should be understood in terms of enabled abilities, but my main reason for this claim (apart from its good fit with the action-space view) is my endorsement of a general view on mental content

⁸⁷ Further considerations in favour of understanding conscious perception in terms of enabled abilities come from the arguments in chapter four, where I suggest that such an understanding allows for a good account of colour perception, and chapter six, where I argue that such an understanding helps us see how to address the explanatory gap.

that space does not allow me to detail here⁸⁸. However, the claim is in any case not mandatory for the action-space view, which might grant that informational states or subpersonal processes should not be understood in terms of enabled abilities, whilst still holding that perceptual experience should. This is because, according to the action-space account, such processes or states contribute to the content and character of a conscious experience only when the personal-level abilities they enable are apt for integration with the agent's ongoing practical reasoning. Even if those states and processes have features that should not be characterised in terms of enabled abilities, the action-space account holds that only the features that should be so characterised are relevant to an account of conscious experience – subpersonal processing and informational states are only relevant to experience insofar as they enable the personal-level abilities to discriminate and classify to which the action-space account appeals.

Our chief motivation for this emphasis on the integrability of perceptually enabled skills thus far⁸⁹ is our consideration of the DVS results in 2.2 and 2.4, which suggested that an important difference between the information in the dorsal stream (which apparently fails to contribute to visual consciousness) and that in the ventral stream (which seems especially implicated in visual consciousness) is the availability of the latter information to practical reasoning and action-selection. The action-space account's gloss on the case of DF illustrates this requirement. As we saw in 2.4, the action-space theorist claims that DF retains spared perceptual sensitivity to aspects of visual form, but that this sensitivity has no upshot in her visual experience since it is not automatically apt to be integrated with her capacities for practical reasoning and action-selection – instead, DF requires a prompt from experimenters, or via inferential strategies that she has learnt herself if her spared sensitivity is to be put to use in pursuit of her current goals.

So, these are the reasons we have seen so far for the view that conscious experience consists in an understanding of the practical consequences of perception – of knowing poise over a space of actions. Before rounding off this chapter by considering some objections to the proposal not already dealt with, let's look at the diagnoses of two other pathologies of consciousness that the action-space account suggests. Firstly, consider blindsight⁹⁰. Blindsight subjects have no awareness of visual stimuli in some or all of their visual field, yet can, under forced-choice conditions make discriminations well above chance

⁸⁸ See Hurley (1998, ch.8) on the debate between informational and teleological theories of content and its relevance to the enactive approach to perception for the sorts of considerations I have in mind.

⁸⁹ Further advantages of this requirement become apparent in chapters 5 and 6.

⁹⁰ See Cowey (2004) for a good review.

concerning colour, orientation, shape, motion and (interestingly) the emotional content of facial expressions. Clearly, subjects retain substantial perceptual sensitivity to certain features, but (like DF) cannot put this sensitivity to use in their current project of a discrimination task (or any relevant projects outwith a laboratory setting) without an experimenter's prompting them to make their discrimination. According to the action-space account, the fact that blindsight subjects typically report no conscious experience of the features on the basis of which they make their discriminations is explained by the fact that the abilities enabled by their perceptual sensitivity to those features are not apt for integration with their capacities for practical reasoning and action-planning.

Secondly, consider unilateral neglect⁹¹. In this condition, subjects fail to report or act upon objects and features in one half of their visual field (typically the left half, after right-hemisphere brain damage) – for example, they might shave only one side of their face, eat food on only one side of their plate, or draw only one half of a perceived object, reporting no awareness of anything unusual about their actions. However, despite their inability to report or act upon information from one half of their visual field, there is evidence that information from that hemi-field is nonetheless being processed to a substantial level. For example, semantic priming effects have been demonstrated where objects in the neglected part of the visual field improve response times to words that are related in meaning to the neglected object – a neglected picture of a tree will facilitate a faster response to the word 'apple' than a neglected picture of a bed (McGlinchey-Berroth (1997)). The action-space account suggests that various abilities are still enabled by the perceptual sensitivity of neglect patients, as evinced by their performance in such tasks, but that these abilities have no upshot in conscious experience insofar as they fail to be apt for integration with the subjects ongoing practical reasoning – the perceptual abilities cannot be automatically used to achieve any of the current goals of the subjects, and so fail to issue in conscious experience.

The foregoing discussion, I think, should be enough to at least motivate the conclusion that an action-space account is worth pursuing. I end the chapter by responding to two sorts of counterexamples to the action-space approach.

3.5 Two Counter-Examples Countered

I have been arguing that conscious experience is a matter of knowledge of the space of actions which one's current perceptual sensitivity to the environment poises one to perform,

⁹¹ See McGlinchey-Berroth (1997) for a review.

of understanding the practical consequences of perception. One sort of counter-example to such an account would be a case where a subject enjoyed a conscious experience, but did not take herself to be able to act upon her environment in the ways that are allegedly characteristic of that experience. This would show that an understanding of the practical consequences of perception is not necessary for conscious experience to obtain. Visual hallucinations and some visual illusions look like plausible cases of this sort. During a visual hallucination there need be no physical object present for the agent to physically engage with, or to sift, sort, classify or compare in any ways the action-space account might wish to emphasise. Similarly, when a subject stares at a Hermann grid and perceives illusory grey dots at the intersections of the white lines, there are no such objects for her to interact with. And, if she is familiar with the illusion, it seems that she will not take herself to be enabled in any ways relating to identifying, tracking or otherwise engaging with grey dots, for she knows that there are none present – her grasp of the space of actions enabled for her will, it seems, include the fact that no actions characteristically appropriate to the presence of fuzzy grey dots are in fact on the cards. In what sense, then, is it the case that such experiences involve grasp of the empowerment of suites of actions, as required by the action-space approach?

According to the action-space account, such instances of illusion and hallucination are cases of misrepresentation. For an agent to be poised over an action-space, and hence for her to undergo a conscious experience (be it veridical or otherwise) is for her to represent the fact that certain possibilities for action are suggested by her perceptible environment, and to do so in such a way that her grasp of those possibilities is apt for integration with her capacities for action-planning and practical reasoning. Instances of illusion and hallucination on this account are cases where some or all of these represented possibilities fail to obtain; where the perceiver's environment fails to conform to the way she implicitly takes it to be. If I hallucinate that there is a pineapple before me, I implicitly take there to be an object that is apt (amongst other things) to be sorted with ovoids and edible things, and sifted from spheres and things to be grasped with a precision grip, when in fact there is none. When I perceive the illusory dots in the Hermann grid, I implicitly take there to be grey dots that occupy roughly such-and-such a set of points in my egocentric space and that are discriminable in shade from both the white of the lines and the black of the squares that surround them, when in fact there are none. Illusions and hallucinations, then, are simply cases in which the agent takes herself to be empowered in ways that she in fact is not⁹².

⁹² Consequently, the 'knowledge' in 'knowing poise over a space of actions', and the 'understanding' in 'understanding of the practical consequences of perception' should not be thought of as factive.

How should this response be squared with cases where the perceiver does not take her experience at face value? When an informed perceiver stares at the Hermann grid, she knows that no actions characteristically appropriate to the presence of fuzzy grey dots are in fact available to her. When a perceiver who knows that she is prone to hallucinate pineapples under certain conditions seems to see a pineapple when those conditions obtain, she might not take her perception at face value, and will thus understand that no pineapple-pertinent actions are currently available to her. Nonetheless, in such cases the action-space account claims that these subjects implicitly take themselves to be empowered to act in ways appropriate to grey dots and pineapples, respectively. Whilst those subjects might *judge* the dots and pineapples to be absent, this is not how they *perceive* them to be – were they to take their perception at face value, they would indeed assent that the actions in terms of which our account wishes to explain these cases were enabled for them. When we look at the Titchener circles or the lines in the Muller-Lyer illusion, they are experienced as being different sizes despite our explicit knowledge to the contrary. In the Titchener circles illusion, picking up one circle looks apt to satisfy goal of picking up the biggest circle in our environment, whereas the other does not – this is something we automatically understand on the basis of our illusory perception of the relative sizes of the circles. It is this level of understanding, rather than the intellectual understanding that is sensitive to the fact that the current appearances are illusory, that the action-space account appeals to in its explanation of conscious experience. The space of actions which an agent implicitly understands to be enabled can come apart from her explicit conceptual judgements about what actions her environment affords. Thus, when our informed perceiver experiences the grey dots at the intersections of the Hermann grid, she can implicitly take herself to be able to sift, sort, track and compare the dots in a certain way whilst explicitly judging that there are no objects present that are appropriate for such actions⁹³.

A second type of counterexample to our account would be a case where an agent takes herself to be poised to act on the environment, and can factor this into her reasoning, planning and intention-forming, but apparently without conscious experience arising. This would show that understanding the practical consequences of perception was not sufficient for conscious experience to obtain. It might be thought that sleepwalkers constitute such

⁹³ In chapter 5 I give a detailed account of how this kind of appeal to practical understanding should be cashed out. We should note here that the practical understanding to which the action-space account appeals is not to be construed as an intellectually demanding achievement, requiring language use, or general conceptual abilities. Were it so, many agents who are intuitively seats of conscious experience (such as certain animals and infants) would fail to be conscious on the action-space view. These issues are also addressed in chapter 5.

cases⁹⁴. Sleepwalkers are capable of navigating their way through an environment and even, in some cases, of performing relatively complex tasks such as driving cars or attempting to carry out mechanical repairs (Cartwright (2004), p.1152). Intuitively, these are examples of agents acting on the basis of their perceptual sensitivity to the actions afforded by the environment, selecting action types and targets appropriately, and apparently acting in a goal-directed manner. If it is correct to describe sleepwalkers as perceptually sensitive to the affordances of their environments and able to put this sensitivity to use in achieving a goal, do they constitute a counterexample to the action-space account?

One option for the action-space theorist is to claim that the sleepwalker does in fact undergo a conscious experience, but is unable to recall that she has done so (Crisp et al (1990) defend such a view). The sleepwalker implicitly understands herself to be poised over a (probably more limited than usual) space of actions, and puts this understanding to use in achieving some goal, thus satisfying the requirements for conscious experience. But due to some inhibition of the systems on which recall and report depend neither she nor we can know about this experience afterwards. Evidence that suggests sleepwalkers are amnesic for a short period after being woken (Cartwright (2004), p.1157) might be taken to support this hypothesis by suggesting that the sleepwalker may have conscious experiences when asleep just as they do shortly after being awoken, but that each such period of conscious experience is unavailable to report and recall.

A second option is to deny both that the sleepwalker has a conscious experience, and that her perceptual situation meets the conditions required by the action-space account. The view that sleepwalkers lack conscious experience perhaps accords best with the popular conception of sleepwalking. It also seems significant that the most commonly cited sleepwalking behaviours such as wandering around, performing a menial household task, and even driving, appear to be behaviours that waking subjects can perform with minimal conscious awareness. The action-space account could perhaps be squared with a denial that sleepwalkers have conscious experience by pointing to discrepancies between the ways in which the perceptual sensitivities of sleepwalkers and normal perceivers to their environments inform their behaviours. The most significant such discrepancy, for our purposes, is that sleepwalking behaviour seems to be inflexibly geared towards the achievement of a single goal, rather than open to the complex and flexible matrix of goals and projects active during waking behaviour. For example, a sleepwalker engaged in cleaning kitchen surfaces might exhibit no sensitivity to the fact that the kitchen is dark, that the surfaces are already clean, or that a concerned family member is asking them

⁹⁴ This has been noted as a problem for action-oriented theories of consciousness in general by Bermudez & Macpherson (1998), para. 32.

what they are doing. This suggests that they are either perceptually insensitive to these facts, or that they are not able to modify their behaviour in the light of such sensitivity, each of which contrasts with the way a conscious perceiver, on the action-space model, must be empowered to act by her environment.

It seems to me, in fact, that the most plausible account of the sleepwalking case lies somewhere between these options. Sleepwalkers present difficult cases for any theory of consciousness, since they manifest some apparent hallmarks of conscious experience (such as using perceptual sensitivity to their environment to inform a goal-directed behaviour) whilst lacking others (such as the abilities to recall and report, and to respond flexibly and intelligently to their environment). As a result, both intuition and empirical studies leave it unclear what we should conclude about the conscious state of the sleepwalker. It seems plausible to conclude that the evidence from both these sources precludes placing sleepwalkers at either end of a conscious/non-conscious continuum. It seems natural to describe sleepwalkers as in a state somewhere between sleep and wakefulness – perhaps we should assume on this basis that their conscious experience has a similarly intermediate status. This looks to be the most natural diagnosis of the sleepwalker's situation, and one that is consistent with the action-space as we have developed it here. The sleepwalker seems to be located somewhere on a continuum between the full and flexible integration of perceptual sensitivity with goals and plans that characterises the normal conscious perceiver, and the kinds of rigid and reflexive responsiveness to the environment that can occur without conscious experience at all. The action-space account thus suggests that the sleepwalker enjoys conscious perceptual experience in proportion to the extent to which her sensitivity to the affordances of the environment can be integrated with her ongoing and long-term goals, wants and plans. Such sensitivity and integrability are both markedly impoverished with respect to normal conscious perceivers, but present to a very limited extent.

Hallucinations, illusions and sleepwalking, then, fail as counterexamples to the claim that conscious experience consists in an understanding of the practical consequences of perception, serving instead to bring out that view in greater detail, and point up some of the explanatory options at its disposal.

Conclusion

We began this chapter by reviewing Pettit's dispositional theory of colour looks, a forerunner of the action-space view. Pettit argues that we should not think of colour experience as involving acquaintance with essentially effect-independent qualia, on the

basis of which we are poised to categorise, discriminate, and act in our environment in various ways. Rather, our perceptual sensitivity to the environment poises us to act in these ways, and this poise is the basis in terms of which we should attempt to explain the content and character of the experience. The discussion of the DVS results in chapter 2, and of Kohler's coloured goggles and synaesthetic colour sensations in this chapter show how this proposal meshes well with various empirical work and puzzle cases. In 3.2 I argued that the sensorimotor theory lacks the resources to explain Kohler's colour-adaptation results, in contrast to the action-space view. In 3.3 I argued for the same conclusion with respect to synaesthetic colour sensations. Each case invites understanding in terms of the range of actions perceptually enabled for a subject, rather than in terms of that subject's knowledge of the likely perceptual consequences of movement.

The second main section of this chapter began by taking stock of our discussion so far. The view that emerges through our consideration of the DVS results, Pettit's theory of colour experience and the cases considered in 3.2 and 3.3 is that conscious experience is a matter of understanding the practical consequences of perceptual sensitivity – of *knowing poise over a space of actions*. On such a model, pathologies of consciousness such as visual form agnosia, blindsight and hemispatial neglect are understood as cases where subjects are unable to put their perceptual sensitivity in touch with their capacities for practical reasoning and action-planning, and thus fail to enjoy conscious experiences that are based on that sensitivity. We ended by looking at the action-space account's responses to a pair of apparent counterexamples. I argued that we should see visual illusions and hallucinations as cases where the subject misunderstands the space of actions currently afforded by her environment, and consequently has an aberrant perceptual experience. Our discussion of these cases brought out the fact that the action-space doesn't construe grasping the practical consequences of perception as an intellectually demanding achievement – thus, the way a subject takes herself to be able to act on her environment can come apart from the explicit conceptual judgements she makes about her environment on the basis of her experience, and consciousness, construed as knowing poise over a space of action, is something that we may share with non-linguistic and non-concept-using agents, such as infants and certain animals. Secondly, I argued that sleepwalking fails as a counterexample to the action-space view, since the sleepwalker is markedly impoverished in the extent to which he can put his perceptual sensitivity to the environment in touch with fluent and intelligent practical reasoning. To the extent that the sleepwalker can do this, they will enjoy conscious experience – a conclusion which looks consistent with the data on sleepwalking.

Our case for the thesis that conscious experience consists in a grasp of the practical consequences of perception is building, but three outstanding issues loom. Firstly, can the action-space account give us a detailed and plausible account of some specific class of experiences, in terms of the sorts of actions and abilities they involve? Secondly, how should we cash out the action-space account's appeal to 'understanding' the practical consequences of movement, and the 'practical reasoning' with which it claims perceptually enabled abilities must be integrated if they are to issue in conscious experience? Thirdly, do we have any reason to believe that the action-space theory can address the sorts of questions about the explanatory gap with which we began in chapter 1? The remaining chapters of this thesis deal with each of these issues in turn. Next, we continue to focus on colour experience as a case study for the action-space view, and develop a theory of the skills and abilities in terms of which such experience should be understood.

Chapter 4: The Agent in Magenta

As we have seen, according to the action-space theory of perception conscious experience is to be explained in terms of an agent's understanding of the space of actions her perceptual sensitivity to the environment currently poises her to carry out. In this chapter, I aim to clarify the action-space proposal by giving a worked-out account of colour perception. Colour experience is possibly philosophy of mind's favourite example of a reductively problematic phenomenal state. Perceiving colour is also, intuitively, a passive phenomenon – colour is a property that 'hits us in the eyes' apparently requiring no effort or activity on our part. If the action-space theory contains the resources for a plausible explanation of the content and character of such experiences, this must count heavily in its favour. I proceed by canvassing the options for both action-space and sensorimotor theories of colour, arguing that each type of theory faces problems that the other can deal with relatively easily, and concluding that we should combine the action-space and sensorimotor views. Crucially, however, I argue that the views on colour perception espoused by sensorimotor theorists in fact suggest only an indirect role for sensorimotor dynamics in colour perception. Thus, the account of colour perception I develop constitutes both another line of evidence that perceptual experience should be accounted for in action-space rather than sensorimotor terms, and a worked-out example of an action-space treatment of a class of perceptual experiences.

4.1 An Action-Space Account of Colour Looks

How might we ground an account of colour perception in enabled perceptual abilities? A natural starting point is to focus on the range of discriminatory abilities perceptual exposure to a colour enables for a perceiver. We have already come across this general idea in Pettit's (2003) account of colour looks as powers of objects to enable a range of sifting, sorting and tracking abilities in a perceiver. One thing that each of Pettit's paradigmatic actions of sifting, sorting and tracking look to have in common is their reliance on the subject's ability to make perceptual discriminations of sameness and difference based on sensitivity to the object's colour. Ability to track an object based on its colour requires an ability to discriminate that object from its background on the basis of its colour, as does the ability to sift an object from a class of differently coloured objects on the basis of its colour. The ability to sort objects into classes of same or similar colours requires the ability to make appropriate discriminations of perceptual similarity.

Following Austen Clark (1993, 2002) and Mohan Matthen (2005, especially ch.11), we can try to ground a theory of colour perception in a catalogue of the range of, and relations between, such discriminatory abilities. Two colours are discriminable from each other for a given perceiver iff the perceiver's same/different judgements in response to the two colours are correct above chance. Knowing the extent of these abilities for a perceiver allows us to chart the space of colours to which they are perceptually sensitive, and the relations of relative similarity and difference which hold between those colours. For example, by presenting a subject with a series of pairs of colour chips, each with marginally different hue properties, but the same level of saturation and brightness, we can move from presenting a subject with a unitary shade of red (a shade which is perceived as containing no traces of yellow or blue) to presenting her with a unitary shade of yellow, whilst at no stage presenting her with a pair of samples she can reliably discriminate from each other. Likewise, we could present her with a series of pairs with a gradual variation in brightness, and move from presenting a pure white chip to a pure black chip in the same way. In fact, for *any* pair of colour samples, we can move from one to the other by a series of presentations of different samples, each one marginally but indiscriminably differently-coloured from the last. In this way, we build up a picture of a multi-dimensional colour-space, its geometry dictated by the extent of a perceiver's sensitivity to similarities and differences between its constituent shades.

Just as we can construct such a colour-space by successive presentation of indiscriminable shades to a subject, we can construct the same space by exploiting the discriminatory abilities the perceiver possesses, by presenting the perceiver with a series of *barely-discriminable* colour samples. Such a space needn't be any less fine-grained than the one constructed via indiscriminable samples – a perceiver may not be able to discriminate red 234 from red 235, or red 235 from red 236 (and so on), but still able to discriminate red 234 from 236, and red 235 from 237 (and so on). Thus, all perceptible shades, and the relations of discriminability between them, are captured in the resultant space. This space is one in which any constituent point can be picked out by specifying its coordinates along three degrees of variation. In the diagram below, an idealised version of the space of colours discriminable to standard human perceivers, the vertical axis represents degree of brightness, the angular coordinate represents degree of hue, and the radial coordinate represents degree of saturation.

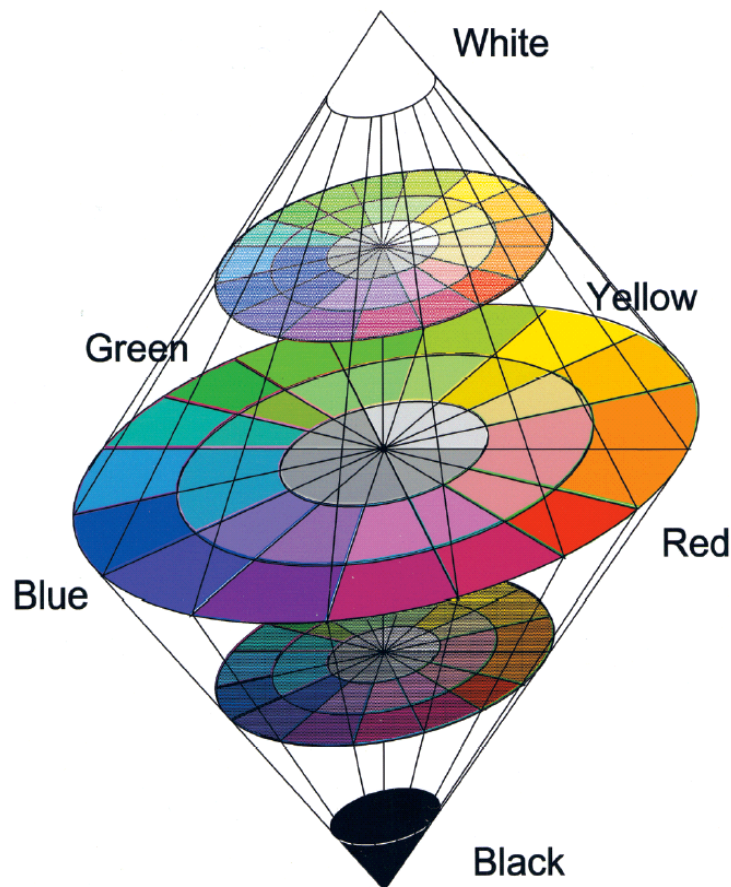


Fig.1

Any point in colour space can be specified by a location in terms of these three coordinates, and the relation between any two points in the space can be given by specifying the vector linking one set of coordinates to the next⁹⁵. The fact that the geometry of this space is dictated by the range of discriminatory abilities possessed by a perceiver gives action a constitutive role in this account of colour perception – our proposal is that the coordinate system which makes up discriminable colour space has enabled discriminatory abilities as its metric⁹⁶. Action-space theorists can use the conceptual apparatus outlined above to claim that an agent's experiences of colour are explained by her implicit knowledge that a certain range of discriminatory abilities are currently enabled for her. The enabling of

⁹⁵ For much more on the precise methods of constructing such a space, including detailed descriptions of the mathematical regularities describing its geometry, and how these are arrived at, see Austen Clark (1993). Matthen (2005), section 2, contains a detailed discussion of sensory ordering in general.

⁹⁶ Glossing such an account as action-oriented requires that we count operations of discrimination, comparison and the like as actions. Recall our defence of this point in 3.4 – firstly, we might think that it makes sense to attribute a discriminatory capacity to a subject only if that capacity can show up in their behaviour, or at least their behavioural dispositions. Secondly, the action-space account claims that discriminatory abilities only have an upshot in perceptual experience when the personal-level abilities they enable are put in touch with capacities for practical reasoning.

these abilities constitutes her occupying a specific point in a complex space of possible enabled discriminatory abilities, the geometry of which is given by the totality of similarity and difference relations obtaining between colours to which the perceiver is sensitive. This is our first pass at an action-space theory of colour perception.

A few points about our proposal. Firstly, since the action space account aims at an explanation of our phenomenology, the discriminatory abilities in terms of which it wishes to explain colour experience had better not presuppose phenomenology, on pain of circularity. But it seems that the discriminatory abilities to which we appealed above don't make such a presupposition – we might suppose that a robot could be successfully programmed to make discriminations between similarly-coloured objects presented to it and output a 'same/different' response without supposing that it experienced the colours it was discriminating. Indeed, as briefly noted in 3.4, some blindsight subjects can make surprisingly fine-grained forced-choice chromatic distinctions between colour-samples presented in their scotoma, whilst reporting that they are not basing their choice on any experience of the colours being presented.

Secondly, note that the set of discriminatory abilities our account will appeal to is rich and complex, as intuitively required of a candidate explanans for colour experience. As noted above, any perceptible colour can be specified by giving its locations on the three axes that structure our colour space. Our proposal is that the metric of those axes be understood in terms of enabled discriminatory abilities – thus, on our proposal, the content of any colour sensation can be captured in terms of a specification of the discriminations currently enabled for a subject. Note also that the range of discriminatory abilities enabled for a subject by exposure to a particular colour contains information not just about the colour currently perceived, but also about all its neighbouring shades and, it seems, all other colours that the subject can perceive. For it turns out that human colour-space has a distinctive and thoroughly asymmetrical structure – for example, red can be more saturated than any other of the unitary hues, red and green can be darker than any other unitary hues, a maximally saturated yellow is closer to an achromatic stimulus than any other maximally saturated colour, and so forth. When we take into account all such peculiarities in the geometry of discriminability space, it becomes apparent that the relations of similarity and difference that obtain between the colours to which we are perceptually sensitive are such that they permit only one type of organisation in discriminability space. Thus we might compare the task of mapping a perceptible colour to a point in colour space to the task of assigning a number to an empty box in a Sudoku grid. The choice of which number to fill in might seem an arbitrary decision, but Sudoku-space

(the grid) has been so designed that for any box, there is one and only one number that can go in it whilst respecting the geometrical rules of the space. Likewise, the discriminability relations that dictate the geometry of colour space permit only one way of mapping a particular colour to a point in that space. This fact means that the enabling of the discriminatory abilities pertaining to a particular colour for a perceiver carries information not just about the relations of similarity and differences obtaining between that colour and its neighbouring shades, but (to some extent) about all the other points in the relationally-defined discriminability space.

These two points might be recruited in responses to perhaps the most obvious objections to our proposal – that the abilities in terms of which we wish to explain colour experience might be enabled for an agent with absent or inverted qualia (recall 1.1). As we just noted, our account entails that the range of abilities the action-space account appeals to are identical in grain and relational structure to the range of colours a perceiver can detect – this might help assuage any absent qualia worries that are motivated by the apparent paucity of enabled abilities as a candidate explanans. However, it still seems legitimate to question why a range of abilities, however complex, should implicate conscious experience – also noted above was the fact that we might design a robot that manifested some kind of sensitivity to chromatic information in its behaviour without any inclination to suppose that it enjoyed colour experience, and that some blindsight subjects seem to possess residual abilities to make chromatic discriminations when prompted, but in the absence of colour experience. We will not be in a position to give the action-space account's response to such worries until the end of chapter 6, which explicitly addresses the relation of the action-space account to the explanatory gap. As a palliative, we can note for now that this problem is not unique to the action-space account, but arises for any attempt to give an account of colour perception in non-phenomenal terms – short of 'explaining' our colour experience in terms of our acquaintance with experiential properties, any theory will have to tackle the difficult question of how the physical properties to which they appeal relate to the experiential properties they wish to explain. The action-space account's response to this question is that the non-phenomenal abilities to which they appeal contribute to the content of a conscious experience only when they are integrated with the subject's capacities for practical reasoning⁹⁷. Thus, the action-space theorist's account of colour blindsight will be an instance of the general account of blindsight given in

⁹⁷ Cf. Austen Clark (2007), who argues that those abilities issue in consciousness only when attention is directed upon them. The action-space account has some affinity with theories of consciousness that appeal to attention in this way (see also Prinz (2000)), since some psychologists define attention in terms of the capacities to put information in touch with practical reasoning to which the action-space account appeals.

3.4: normal perceivers automatically know that the coloured patch in front of them affords matching with such-and-such a range of colours, tracking against such-and-such a range of backgrounds, and so forth – when appropriate, their sensitivity to colour-based affordances is automatically put in touch with their intentional goals and projects, thus issuing in colour experience. This contrasts with the case of the blindsight subject mentioned at the start of this section. This subject retains some of this sensitivity to colour-based affordances, but the intermediary of a prompt from the experimenter is required before this sensitivity can be put to use in the achievement of their current goal (in the context of the experiment) of matching or discriminating. We have been given some empirical reason to think that such integrability of informational states with practical reasoning is key to conscious experience by our discussion of the DVS data in chapter 2, and the diagnoses this hypothesis can provide of visual agnosia, blindsight, and unilateral neglect, and in chapter 6 we will consider conceptual arguments for our view.

Inverted spectrum arguments look easier for the action-space account of colour looks to respond to. It might be objected that a mere appeal to sets of enabled abilities leaves it opaque why the enabling of some such set should look (say) purple, rather than pink. But, as we noted above, it turns out that the structure of our colour-space is thoroughly asymmetrical in a way such that the set of relations of sameness and difference in which one point in colour space stands to the rest is unique to that point. For example, the totality of discriminability relations between colours could not be held fixed whilst only (say) red and yellow were switched around in colour-space, since the action-space account holds that it is constitutive of yellow and red that, when maximally saturated, yellow is more similar to white than to black, and red is more similar to black than to white. Thus, the positions of yellow and red cannot be swapped in colour-space without disrupting the relations of similarity and difference that obtain between black and white and those colours. Were we to attempt to resolve this by swapping the positions of black and white also, we would then need to shift the locations of all the colours constitutively related to black and white by their discriminability relations, and subsequently all the colours constitutively related to *those* colours by their discriminability relations, and so forth⁹⁸ – ultimately, we find that if we attempt to move a point in colour space whilst keeping its relations of discriminability intact, all other points in colour space must move with it.

⁹⁸ In the interests of simplicity, the way I'm stating this example has the misleading implication that each colour stands in constitutive relations of similarity and difference only to certain other colours in that space. But recall our Sudoku comparison, above – the geometry of colour-space is such that every point stands in a unique set of relations to every other point, not just some subset of those points.

And since, on our account, the location of a point in colour space is defined relationally, this is tantamount to that point failing to move at all.

It is worth noting with Austen Clark (1993) and Matthen (2005) that this type of account of colour perception appears to mesh nicely with an empirically well-confirmed proposal concerning the neurological basis of colour-vision -the Hurvich-Jameson opponent-processing model. The proposal is roughly as follows⁹⁹: there are three types of colour-sensitive cones on the retina, each sensitive to different wavelength regions of the visible light spectrum. Clusters of cones feed into three types of colour-coding cells at the output layer of the network – one cell that registers the value of the incoming light on a continuum from red to green, one on a continuum from yellow to blue, and one on a continuum from black to white. Registration works by each cell computing the ratio between the inputs from some combination of cone types (which combination varies according to cell-type), resulting in a measurable degree of excitation/inhibition in each cell. The activation levels of each cell constitute a triplet of outputs that combine to yield a particular value in a three-dimensional space of possible activations.

The interest of this model for our account is that there is considerable evidence that the space of possible activations of the HJ-network is strongly isomorphic to the space of chromatic discriminations available to a perceiver. For example, by attending to the informational sensitivity of a perceiver's colour perception mechanisms using the apparatus of the HJ-network, we can accurately predict for any pair of samples whether or not a subject will be able to reliably discriminate between them. The two main types of colour blindness (impaired ability to make discriminations between red and green in one, or between yellow and blue in the other) are found to correlate with selective impairment of the relevant type of receptor. The fact that our colour vision adapts to systematically distorted light (as when we perceptually adapt to dark or colour-tinted glasses) is predicted by the opponent-process model, which states that prolonged activation of a cell or cluster of cells will result in cell fatigue, with the cell's activation-level tending to return to a neutral value – so consistent exposure to yellow-tinted light will fatigue the yellow/blue cells, and gradually the perceived yellowness of the perceptual environment will wear off. The yellow bias of our chromatic discriminations fades with the activation-levels of the yellow/blue cells. Additionally, the model can reliably predict the colour (and the underlying discriminability-profiles) of our afterimages – in the same way as everything is perceived with a bluish tint after removing yellow sunglasses to

⁹⁹ For a more detailed summary, see Churchland (2005). For original proposal, see Hurvich and Jameson (1957).

which we have perceptually adapted, fixating on a yellow patch then on a plain white background will produce an afterimage of a blue patch. The previous two phenomena are explained by the fatiguing of yellow/blue cells, and the assumption of a correlation between levels HJ-network output levels and the way things appear to the perceiver (the link running via enabled discriminatory abilities, on our account). The assumption that such a correlation obtains yields impressively accurate predictions of the exact chromatic character of our after-images.

The explanatory and predictive power of the opponent-processing model is such that Churchland (2005) has argued that colour sensations should be identified with the output levels of the HJ-network. However, identifying such low-level processing with personal-level sensation seems implausible¹⁰⁰. Surely we should not say that an HJ-network operating in isolation from any other cognitive mechanisms is a seat of colour experience? The action-space account as sketched thus far can make the more modest proposal that the network's output levels enable the discriminatory abilities appealed to in the previous sections, which can go on to form the basis for a colour experience if appropriately integrated into the perceiver's intentional activity. However, as we see in the next section, this approach cannot be the full story about the role of action in colour-perception.

4.2 Problems: The Objectivity of Colour

We can introduce the problems with the above account by noting that it has been set up by appealing to circumstances that differ markedly from those in which we standardly perceive colour. When we perceive colour, it is usually the colour of entire object, displayed under one of an open-ended range of possible lighting conditions, with the object displaying some pattern of shadows and highlights across its surface. It is incredibly rare that we perceive an object coloured and lit so it presents a uniform appearance across its surface. But it seems that such exceptional cases have just been used as the building materials for the theory of colour perception above. To chart the geometry of a perceiver's colour-discriminability space, it was proposed that we investigate which uniform instances of particular shades she can distinguish from uniform instances of others – perhaps by investigating her abilities to discriminate a range of Munsell chips, presented against a uniform, neutral grey background, under fixed lighting conditions. We should wonder whether an account constructed in this manner could possibly do justice to the real, involved and worldly nature of our colour perception.


¹⁰⁰ See Broackes (2003) for a good discussion of the limits of an appeal to opponent-processing mechanisms in an account of colour-perception.

The phenomenon of colour constancy is one way of giving this general worry a more specific expression. The surface of my office wall appears to me as a uniform cream colour. However, due to the patterns of light and shade distributed over it, a patch of the wall in shadow and a patch of the wall in light will enable a very different range of perceptual discriminations in me. Taken in isolation, I might be disposed to sort a patch of the wall in light with white objects, and a patch in deep shadow with black, for example. But I perceive the wall as neither white, nor black, nor (in the most natural sense) as varying in colour across its surface – but as a uniform cream. Likewise, as the sun sets outside my office window and the fluorescent lights overhead become the chief source of illumination, the colours presented by the surface of my desk enable a set of discriminatory abilities in me that differ markedly from those enabled a couple of hours before. However, I have perceived my desk as being identically coloured throughout. The point is not that I *judge* my wall or my desk to be uniformly coloured despite the colour looks they present – rather I *experience* them as uniformly coloured, whilst the fact that they present different colour looks over their surface or passage through time is apparent to me only on reflection¹⁰¹.

The checkerboard illusion also illustrates this worry. In the diagram below, square A looks dark grey and square B looks light grey. But covering up the surrounding squares reveals that they are an identically coloured. Thus, taken in isolation, each square enables an identical range of perceptual discriminations of similarity and difference in the perceiver, whilst appearing very different in colour¹⁰².

¹⁰¹ Thanks to Matt Nudds for pressing this objection against the action-space view.

¹⁰² See <http://www.lottolab.org/Brightness%20illusions%20page.html#> for analogous illusions with different colours, such as orange and brown.




Edward H. Adelson

Fig. 2

Similarly, in the illusion below, the left central cross appears blue, whilst the right appears yellow. Viewed in isolation from their surroundings, however, each is a neutral shade of grey¹⁰³.

¹⁰³ See <http://www.lottolab.org/Brightness%20illusions%20page.html#> for a version of this illusion where the yellow and blue filters can be manually adjusted to strengthen and weaken the effect. The designer of this illusion claims that the central crosses are identically coloured, which I refuse to believe. They are both grey, though.



So, according to the action-space account as developed so far, each cross should enable a similar range of discriminatory abilities in the perceiver, and so look the same. It seems unable to explain how the crosses are experienced as opponent colours, despite this similarity in the abilities their colours enable.

It seems, then, that perceivers can see something as being uniformly coloured whilst it exhibits differing colour-looks, and as differently coloured whilst it exhibits the same colour look. The cases of colour constancy and illusion presented above show that an object's presenting a particular colour-look to a perceiver is neither necessary nor sufficient for that perceiver to experience it as being a particular colour.

A more general worry for the action-space account is that colours, for the most part, are not experienced as subjective, observer-dependent properties, but rather as enduring and objective properties of objects. The ways in which we use colours to categorise, track and re-identify objects over time illustrate this aspect of our conception of colours. But the action-space account of colour has a subjectivist slant – by explaining the colour an object appears to a perceiver by reference to the range of chromatic discriminatory abilities enabled in the perceiver, it makes colours 'response-dependent' properties, whose character is dependent on their effects on the subject who perceives them. This is in tension with the objective aspect of the phenomenology of colour perception.

Can the action-space account accommodate the perceived objectivity of colour? We cannot objectivise colour-looks via an appeal to the physical properties of coloured objects; the physical constitution of an object is compatible with its presenting a wide range of different colour-looks under different conditions of contrast, lighting and perceptual adaptation, so we cannot establish a reliable connection between its physical properties and the colour-look it presents. In light of this, some theorists¹⁰⁴ build a specification of these conditions into their account, holding that colour-looks signal the objective property of being a certain type of physical object, viewed under a certain combination of such conditions. By including a complete specification of the context in which the look occurs, this gives us a link between colour-looks and an objective (though complex) property of being a certain type of physical object, viewed under a certain range of lighting conditions, against a certain sort of background, by a particular perceiver in a certain state of perceptual adaptation. But this doesn't preserve the objectivity of colour experience in the sense we were interested in. Part of the intuition we wished to honour was that colour experience seems to present us with properties of external objects that persist over time, and that can be manifest to various perceivers. This cannot be achieved by indexing the objective property we are acquainted with in colour experience to a single combination of conditions of the object, its surroundings, its lighting and the perceiver.

Matthen (2005) pursues an alternative strategy, proposing that colour-looks acquaint us with the property a coloured-object visually *appears* to have. On this account, since we are only put in touch with the way things appear, the object need not be coloured as it appears to be, nor need the perceiver believe that the object is so coloured. When things are suffused in artificial red light they appear redder than usual without having changed their colour, and without the perceiver necessarily having a tendency to believe them to be reddishly coloured. But the reddish look of the scene still has objective purport since were the perceiver to take her perception at face value, she would take everything around her to be tinted red.

But this proposal neglects what we have just seen - that an object's having a certain colour look is consistent with various different ways of appearing to be coloured to a perceiver (recall the checkerboard illusion), and it's appearing to be coloured in a certain way is consistent with its presenting various different colour-looks (recall colour constancy). Matthen claims that colour-looks put us in touch with the colour which things appear to have, but we have seen that something exhibiting a certain colour-look is neither necessary nor sufficient for its appearing to be coloured in a certain way. The real problem

¹⁰⁴ E.g. Cohen (2000), McLaughlin (2003)

of objectivity facing colour-look-based theories of colour perception is thus that reflection on the way we perceive the enduring colours of objects suggests that the colour-look presented by an object usually fails to determine the colour that object appears. In order to do justice to fundamental facts about colour perception, then, an account must appeal to more than just colour-looks.

4.3 Colours as Ways of Changing the Light: A Sensorimotor Theory of Colour Perception?

The action-space account thus far can give us a plausible account of colour-looks – it can tell us how and why the colours of the two squares in the checkerboard illusion look, in one sense, the same, and how and why the colours of the shadowed and lit portions of my office wall look, in one sense, different. The considerations of the previous two sections showed that such an account leaves out much of importance about colour perception. In the most natural sense, we perceive the squares of the checkerboard illusion as differently coloured, and my office wall as uniformly coloured. Relatedly, an account of colour perception in terms of colour-looks struggles to do justice to the way we perceive colours as objective and enduring properties of objects.

Justin Broackes (1992, 2007) argues for a conception of colours that looks able to deal well with the objections levelled at proposals in terms of colour-looks – that colours are ways in which surfaces change incident into reflected light. He marshals various considerations in favour of this. Think, first of all, of our tendency to look at an object in a variety of lighting conditions and from a variety of angles when trying to establish its colour. A theory of colour perception based on colour-looks might suggest that this behaviour constitutes our trying to find some optimal lighting condition or angle from which the object will exhibit a colour-look corresponding to its true colour. But considerable doubt has been cast on the claim that a colour-look, taken in isolation, serves as a reliable indicator of the objective colour-property a surface possesses. Such a proposal also faces the task of giving a principled specification of the optimal conditions for colour-perception. But is it plausible to think such ideal conditions exist? Why should the colour-look presented by an apple against a backdrop of green leaves in the midday sun be a more objective indicator of its colour than the look it presents against a background of brown earth in early evening¹⁰⁵? Broackes suggests that this behaviour does not aim ultimately at eliciting

¹⁰⁵ See Hardin (1988) for a detailed discussion of the difficulties facing attempts to specify ‘standard conditions’ for colour perception.

some favoured colour-look from the object, but at gaining a detailed appreciation of the constant way in which the object's surface modifies various different types of incident light into reflected light.

Consider also the phenomenon of aspect-shift in colour perception. Broackes (1992, p.460) describes the experience of viewing the cover of a book from a particular viewpoint (one from which the visual information received is ambiguous with regards to the way in which the cover is lit) and experiencing its colour as indeterminate between dark blue and black. At this stage, aspect-shifts in the perception of the book's colour can occur, whereby it is one moment perceived as dark blue, the next as black, sometimes even with the shifts occurring under the intentional control of the perceiver. However, upon further exploration (moving the book with respect to the light-source, or the perceiver with respect to both of these), the book's cover will be perceived as its proper colour, and such aspect-shifts can no longer occur. Similar aspect-shifts can occur when looking at the checkerboard illusion above. This can be explained as due to ambiguities in the light reaching the eye regarding the way in which the object is changing the light – in the former case, it might be a black book in good light, or a dark blue book in weak light. As we have seen, the same colour-look can signal different objective colour properties. Only further exploration can resolve this ambiguity and result in veridical perception of the book's colour.

Perhaps most interestingly, Broackes notes that he himself has impaired red/green colour vision, but he denies that red, green and grey objects present uniform appearances to him or other colour-blind subjects. Grass is experienced as green, old phone-boxes as red, and Broackes is equally good at perceiving the colours of large blobs of red and green paint, ruling out the possibility that his colour experience is grounded in semantic knowledge of the perceived objects. His problem is that he is occasionally prone on a first viewing to take objects to be red or brown when they are, in fact, green. However, when he explores further, an aspect-shift occurs, and the object is perceived as its true colour. Broackes' conclusion is that "The person with red-green deficiencies is simply less good at telling from one viewing what is the object's *way of changing the light*; but by getting a variety of views of it, he may nonetheless recognise that property." (Broackes, 1992, p.462) He gives an analogy with a perceiver with monocular vision: we are not tempted to say of such a perceiver that she experiences the world as two-dimensional. Rather, her visual experience of spatial properties can be ambiguous where that of other people is not. But these ambiguities can be resolved by further exploration. Again, this is a damaging piece of evidence against a version of the above action-space account that relies on the opponent-processing model as the underpinning of the abilities constitutive of colour experience. If perceptions of red and green were constituted by discriminatory abilities underpinned by the red/green channel of the opponent-processing network, we should not expect red/green

colour-blind subjects to experience those colours at all, contrary to Broackes' experience. Even without reliance on the opponent-processing model, it is not clear how the action-space theory as sketched thus far can deal with this fact. For there are undeniably chromatic discriminations of the sort the theory appeals to when constructing a colour-space that are unavailable to Broackes and colour-blind perceivers. However, it is being claimed that colour-blind perceivers can have veridical experiences of the colours that our theory claimed were grounded in the discriminatory abilities that they lack¹⁰⁶.

This view of colour can deal well with the checkerboard illusion and colour constancy cases from the previous section. When we look at the checkerboard, we take square A to be freely illuminated, and square B to be in shadow. Though there is a way in which the two squares share a look, an object can look that way either by being a light grey object freely illuminated, or a dark grey object in shadow. Since we take A to be freely illuminated, and B to be in shadow, we perceive A as light grey and B as dark grey¹⁰⁷. We can provide a similar diagnosis of the way we perceive the coloured crosses in the second illusion. Though both crosses are a similar shade of grey, we perceive the cross on the left as being viewed through a yellow filter, or yellow light, and the cross on the right as viewed through a blue filter, or blue light. Though each cross presents a similar colour-look when viewed in isolation, we implicitly understand that such a look can signal a blue object viewed in yellow light, or a yellow object viewed in blue light. We take the left cross to be viewed through yellow light and the right to be viewed through blue light, and so perceive the left cross as blue and the right as yellow. Similarly, the different colour-looks displayed by different portions of my office wall provide me with information about the constant way the wall modifies different types of incident light into reflected light.

Thus, the colour a surface appears to be is a function not of the colour-look presented by the surface, but by the way we implicitly take it to be changing incident into reflected light. This view of matters can make sense of the objectivist intuition from the previous section – colour perception is perception of a surface's objective property of modifying light in a constant way.

¹⁰⁶ See Broackes (1992, 2007) for further examples and arguments motivating this view.

¹⁰⁷ The effect is no doubt enhanced by our perception of the pattern of the board, and the contrast effects resulting from A's being surrounded by light squares and space, and B's being surrounded by dark squares. However, a range of similar (though slightly less impressive) illusions demonstrate that the effect can be independent of these factors. See: <http://www.lotto1ab.org/Brightness%20illusions%20page.html#>

Just as our colour-look-based theory found empirical support from the H-J opponent-processing model, a conception of colours as ways of changing light finds empirical support in recent work on colour by Philipona and O'Regan (2006) (henceforth POR). POR begin by noting that existing theories of colour perception struggle to explain cross-cultural commonalities in colour-naming and judgements of unique hue – judgements that a colour-sample is a 'pure' instance of a colour, not tinged with other colours. Cross-cultural surveys on colour (Berlin and Kay (1969)) find that there is something special about the colours red, green, blue and yellow: in widely separated cultures, these are the four colours most likely to be given individual names, and the only colours with instances that are judged to be unique hues. These patterns of naming and judging are sufficiently robust across cultures in diverse environments that it is implausible to attempt to explain this fact by appeal to the colours that most commonly occur in perceivers' surroundings.

How could our colour-look-based theory explain these facts? Reflection on the structure of discriminability space leaves it unclear why some points on it should be more privileged than others, either by being accorded a status as a unique hue, or by being likely to be given a name by diverse cultures. It might seem that a theory that makes appeal to the opponent-processing model is well-placed to explain these facts. Recall that according to the opponent-processing model, colours are coded by their values on a red/green axis, and a yellow/blue axis. Perhaps if human colour vision across all cultures proceeds by measuring the light falling on the retina using those colours as parameters, we might reasonably expect to find an echo of this in linguistic practises across cultures. However, a case would need to be made for this response that did not rely on an oversimplified mapping of features of subpersonal processing onto features of experience and linguistic practise. There are, no doubt, many perceptible properties that are processed via measurement against subpersonal parameters that are quite opaque to us, having no echo in experience or linguistic practise. Detection of 'zero-crossings' might be a crucial part of the subpersonal underpinnings of object or scene perception without being a category that shows up cross-culturally in language, or strikes us as experientially salient. It might nonetheless be possible to explain colour-naming practises via an appeal to opponent-processing mechanisms, but further arguments are required to demonstrate this.

Additionally, it is difficult to find an explanation for the cross-cultural similarities in judgements of unique hues in terms of the opponent-processing model. Maximal activation of the red/green channel does not result in perception as of the most reddish red, nor does minimal activation result in perception as of the most greenish green. The same is true of the yellow/blue channel and its corresponding colours.

POR propose a theory that accounts for these features of our colour experience. For each of a set of 1600 Munsell chips (as well as a large array of naturally occurring surfaces), they calculated a function of the chip's reflectance properties, in terms of the way the chip's surface transformed the perceptible incident light (specified by a three-valued vector) into perceptible reflected light (specified by another three-valued vector)¹⁰⁸. This gives us a function for each coloured surface which tells us, given the properties of the light incident upon it, what the properties of the light reflected from it will be. POR found that the functions corresponding to the chips most frequently judged to have a unique hue had what they term "singular reflectance properties" – whilst non-unique hues transform all three values of incident light into different values of reflected light, unique hues transformed the light along two or fewer (as in the cases of black and white) dimensions. In light of this, a 'singularity index' can be constructed, against which the complexity of the way any given colour changes incident to reflected light can be measured. When the colours most likely to be given names, or judged to be unique hues, are measured against this singularity index, it turns out that the hues most likely to be judged unique or named are those that transform incident into reflected light in the simplest way. POR claim that their singularity index can predict, to within the range of a single Munsell chip, which hues a perceiver will judge to be unique, and that the index accurately tracks the probability that a colour will be given a name. The natural conclusion from this data (though one that POR apparently resist drawing in their (2006)) is that perceiving colour involves sensitivity to these reflectance profiles. These results, coupled with the considerations adduced above, make a compelling case for understanding colours as ways of changing light.

The work of both Broackes and POR has been taken as support for a sensorimotor theory of colour perception – a theory which states that the content and character of colour perception is determined by perceivers' expectations concerning how their perceptions will change as a result of certain movements¹⁰⁹. Colour is perceived by grasping the ways in which colour appearances will vary with movements of the perceiver, and changes in other colour-critical conditions – our grasp of the sensorimotor contingencies relevant to colour perception enables us to perceive the constant way an object modifies light, and thus

¹⁰⁸ This differs from the physical notion of a surface spectral reflectance profile, since here the reflectance properties of the objects are identified with respect to incident and reflected light with only three degrees of variation, corresponding to the variations to which the human visual system is sensitive. The surface spectral reflectance of an object is calculated with references to all the differences between incident and reflected light, not just those to which human perceivers are sensitive.

¹⁰⁹ I assume for now that this interpretation is warranted. This assumption will be examined in section 4.5.

its colour¹¹⁰. We have just seen that such a view can deal well with the criticisms of the action-space account considered in the previous section. But we shall now see that just as the action-space theory faced a series of problems related to the objective dimension of colour perception, the theory we have just sketched has pressing problems accounting for its *subjective* dimension.

4.4 Problems: The Subjectivity of Colour

Many of the problems faced by the action-space theory of colour stemmed from its difficulty in accommodating the plausible observation that colours are presented to us in perception as objective properties. Perhaps a more puzzling, but equally plausible observation about colour perception is the fact that the nature of a colour can be at least partially known via transitory and non-standard experiences such as perceiving an after-image, or a neuroscientifically-induced flash of colour¹¹¹. For example, if Jackson's (1982) Mary can induce a momentary illusory perception of a flash of red mist in herself, intuitively this will suffice for her to 'know what it's like' to perceive red, and recognise, remember and imagine objects with that colour. It is difficult to reconcile this observation about the ease with which colours can be known with the last section's sensorimotor account. Implicit expectations about the ways in which appearance will vary with movement do not seem to apply to afterimages or flashes of red mist, and neither experience can be plausibly described as involving taking an object to be modifying light in a certain way. Yet it is clear that each case involves an experience of colour, albeit an aberrant one.

We saw in the previous chapter that the sensorimotor view faces similar problems in accounting for the perception of synaesthetic colours. The problem was that subjects lack sensorimotor expectations about how their photism will change with movement, and as lighting and other perceptually relevant conditions differ. More generally, talk of sensorimotor skills implicated in perceiving a photism again seems inappropriate, since there is no worldly object for the subject to perceptually engage and interact with. Because of the lack of commonality between the sensorimotor profiles of the sensations of veridical and synaesthetic red, the sensorimotor view could not explain their phenomenological similarity. The view of colour in outlined in the previous section has similar difficulties

¹¹⁰See Noë (2004, p.129-132) for a sampling of the range of sensorimotor contingencies he believes a perceiver must grasp to see colour.

¹¹¹ See e.g. Matthen's (2005, p.249-250) discussion of Johnston's "Revelation Thesis", and his own "Transparency Thesis" for further motivation of this observation.

in dealing with synaesthetic sensations, since synaesthetes do not take their photisms to be ways of modifying light, or have implicit expectations about the different colour looks a photism would present in different circumstances.

The view of the previous section also has problems accommodating experiences of 'impossible colours'. For example, fixating for a while on the blue patch below then on the middle of the black square to its right will result in an after-image that appears dark black and simultaneously orange-hued.




Fig.3

The physics of coloured objects precludes there ever being such a colour in the world. Not only do we lack sensorimotor expectations about possible variations in our afterimage, the colour of the afterimage does not correspond to a possible way an object could modify light. The sensorimotor theorist also faces the challenge of specifying what the sensorimotor contingencies involved in the perception of such a colour could be, and explaining how perceivers can have implicit knowledge of these contingencies for colours which they have never before experienced, and would previously have deemed impossible¹¹².

Lastly, a less esoteric worry is whether the sensorimotor account can explain the sense in which the squares in the checkerboard illusion appear the *same* colour, and my office wall appears *differently* coloured across its surface. Though this isn't the way we usually perceive them, it seems undeniable that they sometimes appear this way. Looking at my wall, I can see it as now uniform in colour, now varying – how is this possible if my perception of its colour only involves my implicitly taking it to be modifying light in a certain way?

Perhaps a sensorimotor theorist should respond that the different ways of perceiving the wall correspond to different ways of taking it to be modifying light. One problem with

¹¹²The figure is from Churchland (2005), p.548. See Churchland for samples which produce impossibly dark versions, impossibly bright versions, and impossibly saturated versions of various different hues, and an explanation of this in terms of opponent-processing models. Recall also the affinity between such models and the action-space account noted in section 2.

this line of response is that considering it shows that understanding the sense in which we ‘take’ an object to be changing the light is problematic. If such ‘taking’ is supposed to be something that *I know about*, then to the extent that it is plausible to suppose I possess such knowledge, I know the wall to be modifying light in a constant way across its surface, not in a way that differs according to whether a portion is in light or shade. On the other hand, if the way the wall modifies light is supposed to be something that I automatically grasp in perception, how can we explain my ability to attend selectively to the different ways of perceiving it?¹¹³

Just as the sensorimotor account could readily meet the challenges posed for the action-space account, the action-space account is ideally placed to accommodate the cases considered in this section, each of which poses a problem for, if not a counterexample to, the sensorimotor theory. This is clearest for the cases of colour constancy and the checkerboard illusion just considered – recall from section 3 that the action-space account tells us why the portions of my office wall look differently coloured, but not why they also appear uniformly coloured. We can see how it can handle the other cases from this section by recalling its treatment of synaesthetic colours (3.3). I argued that in contrast to the lack of commonalities between the sensorimotor profiles of synaesthetic and worldly colours, there are substantial similarities in the suites of actions empowered in each type of colour experience, such as pop-out, priming, interference and Stroop effects. Ultimately, the perceptual similarities between synaesthetic and worldly colours are explained by such similarities in the way a subject is poised to sift, sort and track those colours in a way that respects their perceptible similarities and dissimilarities to other colours to which the subject is perceptually sensitive.

The action-space account treats after-images analogously. As with synaesthetic colours, the after-images a perceiver can experience can be located in a multi-dimensional space whose geometry is dictated by the range of discriminations the perceiver can and cannot reliably make between its constituent points. After-images are experienced as similar to worldly colours insofar as they stand in relations of relative similarity and difference to other perceptible after-images that are similar to the discriminability relations that obtain between worldly perceptible colours. As we have seen, the space of perceptible after-images will include colours that could not exist in reality, such as yellowish blacks. The content and character of such impossible colours is likewise explained by the

¹¹³ Moreover, we might question whether a response suggesting that we switch between two ways of perceiving the wall is phenomenologically apt – when we have attended to both of them, don’t we see both the constancy and the variation in its colour at once? See Noë (forthcoming) for discussion of this topic.

discriminability relations they stand in to the other colours, standard and impossible, that the perceiver can experience. Were it possible to make Munsell chips yellowish-black, and other impossible colours, these colours would have been charted by the procedure outlined in section 2. The fact that such afterimages are both *impossible* and perceptibly *coloured* shows that the space of discriminations available to a subject outstrips the space of physically possible colours¹¹⁴.

In light of all this, we can see that the action-space account can readily accommodate the point about the subjectivity of colour with which we began this section. According to the action-space view, an after-image or neuroscientifically-induced flash suffices to give a perceiver knowledge of a colour since the perceiver automatically and implicitly grasps both its place in a network of similar and different colours, and the discriminatory and other abilities that are enabled in virtue of this.

4.5 A Two-Level Enactive Theory of Colour Perception

We have considered two types of enactive theory of colour perception. Each enjoys strong empirical support, and can give promising treatments of cases that confound the other view. But each also faces serious problems from cases and observations that the other can easily handle. It is natural to wonder whether this situation could be resolved by combining the two views.

The two views are *prima facie* compatible, and there is even reason to think that the sensorimotor view stands in independent need of something like the action-space view as a component. The sensorimotor theorist claims that we perceive colours by understanding a set of possible ways their appearances could change according to different viewing conditions. But the sensorimotor theory as it stands gives no account of these changing appearances, only of the higher level property explained by perceivers' grasp of the structures of such possible changes. In response, we can hold that the action-space theory provides the necessary explanation of colour appearances.

The sensorimotor theorist might respond that no such independent account is required - the colour appearances they appeal to can be accounted for in terms of the systematic relations obtaining between the sensorimotor profiles of colours, rather than in terms of some more

¹¹⁴ Recall that a neat explanation of the extent of these abilities is afforded by taking them to be partly enabled by the H-J opponent processing network (Churchland 2005)

basic account. On such a view, a perceiver understands that an object's look would differ a certain way in different viewing conditions by understanding the different sensorimotor profile that the object would look to have in such viewing conditions. But this will not work. The major advantage of the sensorimotor theory was that, by identifying an object's colour with the objective property of modifying light in a constant way, it could accommodate intuitions and cases that turn on the objective aspect of colour perception. But the solution under consideration shifts from viewing colour perception as a matter of a perceiver latching on to an objective property via their grasp of sensorimotor relations, to viewing it as a matter of a perceiver taking an object to look as if it has a certain sensorimotor profile, that stands in systematic relations to other sensorimotor profiles the object would look to have in different circumstances. This undermines the objective aspect that made the sensorimotor theory attractive. If the objectivity of the theory is to be maintained, grasping a sensorimotor profile cannot be accounted for in terms of understanding what closely related sensorimotor profiles would look like, since this makes grasp of sensorimotor profiles holistically defined in terms of a perceiver's implicit knowledge of the profiles' relations of apparent similarity and difference in just the way that resulted in a problematically subjectivist account for the action-space theorist. To preserve the objectivity of the account, then, the sensorimotor profile an object possesses cannot be explained in terms of how it and related profiles appear to the perceiver - an object's having a certain sensorimotor profile must be an objective matter. We might appeal to subjective considerations, such as understanding of the systematic relationships between actual and possible patterns of looks, to explain how a perceiver comes to be able to grasp the objective sensorimotor profile, but such an understanding cannot be what *constitutes* this grasp, on pain of the sensorimotor account surrendering its objective aspect¹¹⁵. This point about objectivity, coupled with the sensorimotor theory's inability to give an account of cases with which the action-space account can deal naturally, shows that we need an account of the underpinnings of the changing appearances to which the sensorimotor theory appeals, and the arguments of sections 4.1 and 4.4 give us good reasons to think that this account should be in action-space terms. But the point also gives us reason to question whether we were right to think of the work of Broackes and POR as motivating a sensorimotor view in the first place.

¹¹⁵ A simpler though less decisive way of arguing against this sensorimotor solution would be an appeal to the subjectivist intuition that there must be some appearances that stand behind sensorimotor profiles and need explained – for example, that we can remain neutral on how we take something to be lit, and still think that it looks a certain way. If we are convinced that such a look needs an explanation, we will need to look past the sensorimotor view to find one.

Our point was that to accommodate objectivist facts about colour, the property to which we appeal must be an objective property of the object, not merely a function of the way things appear to the perceiver. The conception of colours as ways of modifying light is attractive since it provides a well-motivated way of honouring this point. But how is identifying colour with this objective property supposed to support the sensorimotor view? What motivates the sensorimotor theorists' claim that POR's work uncovers "the intrinsic sensorimotor structure of colours" (Hurley and Noë (2007, p.10))? Their intuitive idea is that knowledge of the way colour appearances change with movement and other conditions is necessary to grasp the objective colour property an object has. But compare this way of coming to perceiver an object's colour with the way we come to perceive its shape. The shape of an object can be revealed to me by my understanding of the ways in which its appearance changes as I move around it. But this does not mean we should have a sensorimotor theory of what it is for an object to have a certain shape, or conclude that the objective property of shape has an 'intrinsic sensorimotor structure'. Rather, shape is an objective property that can be grasped in perception *as a result of* sensorimotor knowledge. Likewise, colour is an objective property of objects, perceptual grasp of which can be enabled by sensorimotor knowledge.

When considering the support afforded to the sensorimotor view by the work of Broackes and POR, we should also note that a perceiver's grasp of the possible perceptual effects of their movement with respect to an object is, on their views, only part of what might be involved in coming to see how that object modifies light. Understanding the different ways that the object might look under different sorts of light source, when in the shade of another object, or as other colour-critical conditions vary can all contribute to coming to perceive that an object modifies light in a certain way. And we have been given no compelling reason to think that such relations between appearance and perceptual context should be construed as part of the sensorimotor profile of the object – they look like facts about relations between objects, light sources, and appearances, not about relations between a perceiver's possible movements and appearances. It seems fair to say that some set of counterfactuals contributes to a sensorimotor profile only when they concern the ways in which perceptual appearances change with movement relative to the perceiver – but if we define sensorimotor profiles in this natural way, counterfactuals about variations of appearance with changes in light sources are ineligible for admission, since those changes can obtain without a change in the relative position of the perceiver and the object. We have, then, two reasons for suspicion of the claim that a conception of colours as ways of changing light supports the sensorimotor theory. Firstly, since sensorimotor properties or profiles are relations between possible movements and possible appearances, such

properties or profiles cannot accommodate the point that colour is an objective and enduring properties that made it appealing to see colours as ways of changing light. Secondly, on this view of the nature of sensorimotor properties and profiles, there are many factors that can contribute to a grasp of an object's way of modifying light that should not be understood in sensorimotor terms, such as the relationships between changes in the light falling on a static object, and the way those changes effect the appearances that object presents to a static perceiver.

If this is right, then an understanding of sensorimotor relations can be invoked to explain how grasp of an objective property is *enabled*, but not how it is *constituted*, on pain of losing the objectivity of colour perception. What, then, is constitutive of the perceptual grasp that an object has a certain colour property? I propose that just as perceptual grasp that an object displays a certain colour look consists in the perceiver's implicit knowledge of the range of abilities her perceptual sensitivity enables in her (4.1), perceiving that an object has a certain objective colour-property consists in the perceiver's implicit knowledge of abilities to co-classify, re-identify, and track the object on the basis of her perceptual sensitivity to that property. If this is a tenable option for perceptual grasp of appearance properties, it can be applied to our grasp of higher-level objective colour properties too. Perceiving objective colour properties is, as Broackes and POR suggest, a matter of perceiving the way light is being modified. Knowledge of sensorimotor dynamics – for instance, of the ways in which the pattern of colour-looks displayed by an object or surface would alter with colour-critical conditions – is likely to form part of an explanation of how subjects can grasp this property in perception. But whilst sensorimotor knowledge might help to *enable* such a grasp, I am claiming that the grasp itself *consists in* a subject's ability to act with respect to coloured objects – sifting, sorting, tracking, comparing and re-identifying them – in a way that respects their possession of the property of modifying light in a constant way. Thus, colour perception emerges as another case where appeal to sensorimotor dynamics may play a role, but only in enabling the subject to take herself to be poised over a space of enabled actions in a certain way. Such knowing poise is the essential feature of colour experience, not the sensorimotor dynamics which help to enable it.

The argument of this section has been that the theories of colour-perception of 4.1 and 4.3 need to be combined. But thinking about how the sensorimotor approach could accommodate the objective aspect of colour perception that was supposed to be its strength led us to reconsider the extent to which the conception of colour the sensorimotor theorist appeals to really supports their view. We concluded that understanding sensorimotor dynamics can explain how our sensitivity to objective colours is enabled, but not how it is

constituted, on pain of losing the objectivity of colour perception. In light of this, I proposed that we should extend our action-space account to perception of objective colours. We arrive at a two-level enactive theory of colour perception. At each level, colour perception is a matter of the enabling of a characteristic suite of actions of sifting, co-classification, re-identification, and the like. Understanding of sensorimotor dynamics plays an important role in showing us how the two tiers of our account are related – the suite of classificatory and discriminatory actions constitutive of objective colour perception at our second level is enabled by our grasp of the patterns of and relationships between the colour looks at our first level. Our two-level theory of colour-perception is empirically well-supported and can accommodate both the subjective and objective aspects of colour perception. Whilst each of its two-levels is explained in terms of the action-space theory, it can nonetheless accommodate the important but indirect role that grasp of sensorimotor relations plays in colour perception.

4.6 Synaesthetic Sensations and Special Spectacles Reassessed

We have arrived at a view according to which colour perception has both a subjective and an objective aspect. Most of the time, when we perceive colour we perceive it as an objective and enduring property, and one that an object can manifestly possess by appearing in one of several different ways, depending on lighting and contrast effects, the perceiver's current state of perceptual adaptation, and so forth. We have good reason to think that this property is the way the object modifies incident into reflected light. According to our action-space account, sensitivity to this property contributes to the character of a subject's experience when that subject understands herself to be empowered to act in various ways – sifting, sorting, classifying, comparing, re-identifying – on the basis of that sensitivity. But there is also a subjective aspect to colour perception on our account. This is because we see the colour properties an object possesses *by* or *through* seeing patterns in the ways that the object appears to be coloured, where 'appearing to be coloured' in a certain way is compatible with possessing any one of a variety of objective colour properties. The action-space account of such colour appearance properties follows Austen Clark (1993, 2007) and Pettit (2003) in grounding our experience of such appearances in enabled abilities to sift, sort, track and classify. In contrast to the case where a subject takes herself to be able to act with respect to an object in ways appropriate to its objective colour, a subject may take herself to be empowered to act in ways appropriate to an object's colour *appearance* whilst understanding that a very different set of actions are appropriate with respect to its objective colour. A stark instance of such a case would be when a synaesthete takes herself

to be able to sift, sort and track a grapheme in ways appropriate to green objects, whilst nonetheless understanding that the grapheme should be sorted with other black characters, is more similar in colour to brown than to blue, and so forth. Another such case would be when a subject sees that a spot catching the light on a dark blue vase looks almost white, whilst still understanding that the vase is coloured uniformly all over, and is more similar in colour to black than to white. Coming to see an object as coloured in a certain way involves coming to understand the regularities and significances in the patterns of appearances it presents. In some cases, this might involve having sensorimotor expectancies about the ways the objects appearance would differ with movement. In some cases, it might not – perhaps in some cases the perceiver’s grasp of how an object would look were it lit differently, or were the perceiver in a different state of adaptation suffices for perceiving its objective colour¹¹⁶. Ultimately, however, sensorimotor knowledge is of only instrumental importance for a perceiver’s grasp of the significance of appearance properties – important only to the extent that it contributes to their grasp of the way that an object can be sifted, sorted, tracked and classified according to the way in which it changes the light.

Our account brings out the fact that our previous survey of synaesthetic sensations and special spectacles suppressed certain subtleties. The fact that one phenomenon is dealt with by the subjective, and one by the objective, component of our account of colour shows that the action-space account does not treat them in exactly the same way, a point which we were not in a position to see in the last chapter. But the asymmetry in the action-space account’s treatment of these two cases is well-motivated by our two-level theory, and corresponds to an asymmetry in the way colours are experienced – as subjective phenomena in one case, and as objective properties in the other. This view of the cases suggests an answer to the general question posed by Hurley and Noë (2007) – why don’t synaesthetic colours adapt away? More generally, how can a single theory of colour perception explain both the persistence of synaesthetic colours, and the adaptation to Kohler’s goggles? The action-space account’s response is that the peculiarity of each of these cases lies in their pulling apart one aspect of colour experience from the other that usually complements it. In synaesthetic colour experience, subjective colour looks are divorced from their usual role of indicating, as part of a larger pattern of colour looks, objective colour properties. In

¹¹⁶ Noë (2004, p.169) claims that the way an object changes its appearance changes as a result of the way it’s lit is one of the sensorimotor contingencies relevant to colour perception. However, since movements of the perceiver, object or light source are not necessary for an object to change its appearance as a result of the way it is lit, it seems implausible to suggest that we should understand the relationship between appearance and lighting conditions in terms of patterns of dependence between sensation and movement.

Kohler's adaptation results, we find that perceptual grasp of objective colour properties can, given sufficient time, adapt to radical distortions in the usual patterns of looks underlying those properties. Thus, our account suggests, synaesthetic colours fail to adapt away because synaesthetes do not take them to be indicative of objective colour properties in the same way as standard subjective colour looks – perhaps due to the unusual stimuli those colours attach themselves to, perhaps due to the fact that they fail to figure in larger patterns of subjective looks or display sensitivity to changes in colour-critical conditions in the same way as standard subjective colour looks, or perhaps due to some other feature of the way synaesthetes process those colours, which further understanding of the neural basis of synaesthesia may uncover. By contrast, the distortion in the perception of objective colour properties caused by Kohler's goggles *does* adapt away, since over the course of adaptation perceivers come to manifest an automatic perceptual grasp of a fact that they explicitly knew all along – that the aberrant looks caused by donning the goggles do not reflect the way objective colour properties are distributed. Adaptation consists in the perceiver's acclimatising to the new patterns of subjective looks that are now indicative of particular colour properties, and thus coming to see objective colour properties as distributed in the way they did before donning the goggles. When the goggles are removed, they must relearn the old sets of relations between subjective appearances and objective properties, and thus go through another period of perceptual distortion and eventual adaptation¹¹⁷.

This account of adaptation to Kohler's goggles seems to imply an account of the subject's phenomenology during adaptation process slightly at odds to the one suggested thus far. Standardly, interpreters of Kohler's results speak as if colour vision returns entirely to normal over the course of adaptation (Hurley (1998), Pettit (2003), Hurley and Noë (2007)), and whilst we have not explicitly endorsed this assumption in our discussion of these results, neither have we questioned it. However, if the action-space account claims that adaptation consists in learning to see objects as modifying light in a certain way *through* seeing the aberrant patterns of colour looks introduced by the goggles, we should

¹¹⁷ The account of adaptation to Kohler's goggles I give here looks similar to those provided by Pettit (2003) and Hurley and Noë (2007), each of which invokes sensorimotor expectations on the part of the subject. But any appeal to sensorimotor dynamics is again, on my account, of only indirect relevance to explaining the content and character of perception – Hurley and Noë suggest that adaptation is triggered since subjects implicitly know all along that the colours of objects do not change in response to head and eye-movements as they seem to upon first donning the goggles. On the action-space account, however, subjects only have such sensorimotor expectations because they implicitly know that objects are invariantly apt to be sifted, sorted and categorised according to their colours in ways unaffected by head and eye movements. It is thus implicit knowledge about objects and the ways in which they should be treated, rather than about the ways in which sensation changes with movement, that explains the fact that adaptation occurs.

expect subjects in the experiment to report that colour properties appear to be distributed as they did before the goggles were donned, whilst the patterns of colour looks presented by the world through the goggles remain distorted. Does this prediction – that objects appear to be veridically coloured whilst the subject’s colour sensations remain distorted – even make sense? Is there any evidence in its favour in Kohler’s reports? In fact, strange as our prediction may sound, it seems to square very well with how subjects describe their experiences of adaptation. For example, even after 60 days, at the height of his adaptation, and whilst reporting that objects appear veridically coloured, Kohler says that:

“The distracting effect which the spectacles had in the beginning is completely gone now. Now I actually feel comforted by the sight of a table which looks blue on one side and yellow on the other. I am so used to it that I would get upset if it were otherwise...” (Kohler (1964, p.113)

This suggests that whilst he perceives objects to be coloured as they were before he donned the spectacles, a focus on the way things subjectively appear to him reveals that the distortion in colour looks caused by the spectacles is still present. Similarly, twenty days after the spectacles had been removed (with full re-adaptation occurring after thirty days), Kohler reports:

‘I have become totally immune to the distracting influence of all these novel impressions. I can now work for hours on end without being the least bit inconvenienced by this great variety of discolourations. However, I am still aware that they are there.’ (Ibid. p.115)

This again suggests that adaptation consists in learning to see objective colour properties *through* the distortions in colour appearances that have been put in place, not necessarily in an alteration of those appearances themselves. These kinds of reports are just what we would expect were the two-level view defended here correct.

We also noted in 3.2 that subjective reports of perceptual adaptation are supported by improved performance at a task where the subject must manually adjust the hue displayed by a colour wheel until it contains no traces of blue or yellow. However, we did not mention there that subjects’ performance at this task improves by only fifty percent, a finding

apparently at odds with their reports about the extent of their adaptation¹¹⁸. But when we note that the matching task consists in categorising a single colour patch, under fixed lighting conditions, against a uniform background, and recall the similarity between this and the procedures we outlined to chart the similarity space of colour looks in 4.1, the fact that subjects show only limited adaptation in this context looks less surprising for our two-level view. This particular task is suited to probe the way things subjectively appear to perceivers, and the differences between this context and colour vision in a naturally lit, inhomogeneous and freely-explorable environment make it, on our account, an imperfect measure of how perceivers take objects to be coloured. If subjects adapt with respect to the objective colour properties they perceive things to have, but not with respect to the colour-looks they experience, we should expect such a measure to show imperfect adaptation at best¹¹⁹. Also in keeping with our interpretation here is the fact that subjects' colour categorisations in 'free vision', with none of the restrictions characteristic of the above laboratory set-up adapted almost completely, in contrast to the level of adaptation suggested by the colour-wheel experiment (ibid, p.45). Kohler's adaptation results seem to support our two-level view's division between subjective colour sensations and the way objects in fact appear to be coloured.

We should conclude, then, that our two-level account of colour perception provides a natural way to make sense of the lack of adaptation in synaesthesia, and the presence of adaptation in Kohler's results. It also seems apt to explain the apparent discrepancies in subject's reports, and different measures of adaptation – subjects come to perceive objective colour properties as veridically distributed, whilst their subjective experiences of colour remain distorted. Colour perception, on our account, has a subjective and an objective aspect, but both are to be understood in terms of the ways a subject takes herself to be empowered to act on the basis of her perceptual sensitivity.

Conclusion

¹¹⁸ "A blue of one half its initial intensity on the colour wheel was sufficient to compensate for the yellow of the spectacles, and the blue of the spectacles was judged to be completely equal to gray – in other words, the subject no longer perceived it as colour." (Ibid. p.114) This fact is also neglected in the other enactivist discussions of Kohler's results.

¹¹⁹ There are various possible explanations of why this measure reveals even a fifty percent adaptation available to the action-space theorist. The most simple explanation would be that the measure probes partly for how things are subjectively coloured for the subject, and partly for the objective colour the subject takes things to have. A more complex proposal might be that the way colours subjectively appear and the way things appear to be objectively coloured are not independent of each other – thus, as a subject's objective colour-categorisations adapt to distorted subjective impressions, the subjective impressions might adapt with them to some extent. I won't attempt to adjudicate between these possibilities here.

I have been arguing for a two-level action-space view of colour perception. Perceiving objective colour is a matter of implicitly grasping how objects can be sifted, sorted, tracked and re-identified on the basis of the way in which they modify light. This is enabled by a grasp of the significance of the current patterns of colour-appearances, and how those appearances might vary with colour-critical conditions, and those appearances are in turn explained in terms of a perceiver's abilities to sift, sort and track colours in a way that places them in a set of relations of similarity and difference that make up perceptual colour space. We have seen that this two-level view can accommodate both the subjective and the objective dimension of colour experience, and can deal with problem cases such as colour constancy, colour illusions, afterimages, impossible colours and synaesthetic sensations. Additionally, we saw in the last section that our two-level view accurately predicts the unusual phenomenology and apparent discrepancies between different measures of perceptual adaptation that we find in Kohler's experiments with colour-distorting goggles. A question that arises for such a view, however, is why it deserves to be called a theory of colour *experience* – why not, for example, take it as a thesis about the action-oriented *contents* of colour experiences? In 4.1 and 4.5 I stated that, according to the action-space account, the enabled abilities that explain our colour perception have an upshot in conscious experience of colour only when the subject *understands* that those abilities are enabled for her – when she is able to integrate those abilities with her ongoing practical reasoning. We have been given some reason to endorse this claim by our interpretation of the DVS results in chapter 2, according to which integrability of information with practical reasoning is a functional hallmark of visual consciousness, and perhaps from the plausible treatment the action-space account affords of various cases of blindsight, neglect, and the putative counterexamples considered in 3.5. But what does such a requirement of understanding or integrability really amount to? And how does it put us in a position to meet the challenge from the explanatory gap set out in chapter 1? The next two chapters address these questions – in chapter 5 I say more about the sense in which a perceiver must 'understand' the abilities enabled by her perceptual sensitivity to her environment if they are to issue in conscious experience, and in chapter 6 I discuss how the action-space account at which we have arrived contains the resources for a dissolution of the explanatory gap.

Chapter 5: On Thinking Before You Act

The action-space account we have been developing claims that conscious experience is a matter of implicitly knowing, or having a certain kind of access to, the ways one is poised to act in and upon one's environment. We have seen various empirically-motivated arguments for this proposal, and in the next chapter I shall argue that adopting the action-space framework suggests a promising way of addressing the 'Hard problem' of consciousness with which we began in chapter 1. In this chapter I want to address a lacuna in our sketch of the action-space account thus far – the question of how to cash out the crucial appeal to 'implicit knowledge' of, or access to, or understanding of, enabled actions. I argue that such an appeal presents an apparent problem for enactive theories of consciousness, both action-space and sensorimotor. One advantage of enactive theories was supposed to be their ability to explain the presence of consciousness in creatures without linguistic or conceptual abilities. But the most natural construal of enactive theories' appeals to knowledge or access tie them to just such cognitively demanding abilities, threatening to deny conscious experience to agents who intuitively possess it. Section one brings out the problem, and the constraints on a solution that our characterisation of the problem imposes. Section two outlines Susan Hurley's 'shared circuits model', which will frame our discussion of the solution, and section three clarifies how this bears on the problem as sketched in section one. Section four sums up how the action-space requirement should cash out its appeal to 'implicit knowledge' or 'access' in light of the previous sections. The final sections tie up loose ends, clarifying how the material of this chapter relates to our account of colour experience in section five, and exploring the extent to which our solution suggests a blurring of the boundaries between sensorimotor and action-space accounts of conscious perception in section six.

5.1 The Problem

Recall the reason why sensorimotor and action-space theorists appeal to implicit knowledge or non-inferential access in the first place. If a theory of consciousness is to be grounded in an appeal to sensorimotor dependencies or the actions enabled for an agent, then it must appeal to more than the fact that we can describe a system in terms of the contingencies that obtain between its actions and its perceptions, or the ways in which its sensitivity to its environment enables it to behave. Merely carrying information about the dependencies obtaining between

action and perception cannot suffice for conscious experience. After all, a standard thermostat can do that, as evinced by its ability to respond to an input temperature (perception) with an output signal that will successfully bring about the desired temperature (action). We should reject a theory that makes a standard thermostat a seat of perceptual consciousness. Likewise, it cannot be enough to be perceptually sensitive to the ways an environment affords satisfaction of your goals – after all, our thermostat can do that too, as evinced by its ability to use information about the environmental temperature to achieve its goal of heating the room to a certain level. In light of this, we need to add an extra requirement that perceivers must meet if their enactive credentials are to result in consciousness. Both action-space and sensorimotor theorists have proposed that the agents enactive knowledge must be somehow integrated with thought, planning and reasoning to result in conscious experience. I'll call these '*integration requirements*'.

For example, we find sensorimotor theorists claiming that:

“For a creature (or a machine for that matter) to possess [conscious] visual awareness, what is required is that, in addition to exercising the mastery of the relevant sensorimotor contingencies, it must make use of this exercise for the purposes of thought and planning...To see is to explore one's environment in a way that is mediated by one's mastery of sensorimotor contingencies, *and* to be making use of this mastery in one's planning, reasoning and speech behaviour.”

O'Regan and Noë (2001), p.944

Likewise, the action-space account appeals to implicit knowledge of, or a kind of automatic access to, the space of enabled actions over which the agent is currently poised. Recall that DF (or DF's brain) can use visual information to correctly orient her hand to post a card through a slot, but without conscious experience of the slot's orientation. The action-space account explains this lack of experience in terms of DF's lack of knowledge of the space of actions her perceptual sensitivity to the environment currently affords. The closest we have come to making this claim about her lack of knowledge explicit is suggesting it should be understood in terms of DF's lack of access to the information about the slot's orientation – though this information is clearly being processed (since her behaviour is differentially sensitive to the slot's orientation), DF is unable to factor this information in to her ongoing practical reasoning in the way that a normal perceiver faced with such a task could. Hence

she requires the intermediary of a prompt from the experimenter before this information can be put to use in her current project of the posting task.

Each enactive theory clearly needs some story about why the sensorimotor contingencies and sensitivity to possible actions in terms of which we can characterize the behaviour of a thermostat do not issue in conscious experience. But the brief stories about the availability of the contingencies or enabled actions to reasoning and planning offered thus far risk moving the bar for conscious experience too far in the opposite direction, denying conscious experience to agents who intuitively possess it. For the orthodox way of understanding the capacities for thinking, reasoning and planning to which each account appeals is in terms of a conceptual, linguistic model. The brief way of stating the problem with which I am concerned in this chapter is that the action-space account appeals to an agent's nonconceptual grasp of the space of actions her environment affords, but cashes that grasp out in terms of the integrability of those affordances with capacities for thought, reasoning and planning that are (prima facie) paradigmatically conceptual¹²⁰.

To bring out the less-brief way of stating the problem, we need to sample some of the reasons for viewing thought, planning and reasoning in this way. First, note that it is indeed natural to think of the integration requirements offered by each theory as appeals to the ability to entertain certain (yet to be fully specified) sort of thoughts about sensorimotor contingencies or enabled actions. Each requirement claims that the creature's enactive knowledge must be evinced not merely in its behaviour, but also in its thought, by way of its ability to use this knowledge to inform reasoning and planning. But, for better or worse, the default method employed by philosophers for understanding capacities for thought in general, and reasoning and planning in particular, is a linguistic one, tying it to the ability to token mental states with conceptual content and freely recombinable constituents. I'll call such a view *'the orthodox*

¹²⁰ Another motivation for the integration requirement might be Evans' claim that perceptual information can only become perceptual *experience* "when it serves as the input to a *thinking, concept-applying and reasoning system*; so that the subject's thoughts, plans and deliberations are also systematically dependent on the informational properties of the input. When this happens, we can say that the person, rather than some part of his brain, receives and processes the information." (Evans (1982): 158 original italics). One reason for holding such a view would be the thought that it is a necessary condition on experiences that they be had from a particular perspective. I say a little about how the action-space account can accommodate such a view in the last section of this chapter, on the relation between my position and Hurley's.

conception'. I won't try to do justice here to all the philosophically complex motivations for such a conception, but here is a sample.

It's widely accepted that an essential feature of thoughts is their capacity to stand in certain rational or inferential relations to one another – entertaining one thought can involve another thought's being entailed, justified or ruled out. We might think this because we need our analysis of thought to play a certain role in an epistemological project, or because we think this is a conceptually necessary condition on what it is for something to be a thought, or for both these reasons. Such a view of thought thus owes a story about how one mental state can rationally relate to another, and the standard one is in terms of inferential relations obtaining between those states. Thus, one contentful state can justify another by being able to figure as a premise in a valid argument whose conclusion is the other state. But if we subscribe to such a position, the states to be rationally related must be subject to some principles of decomposability and recomposability – they and their constituents must have a grammar that respects the rational relations in which they stand to other states¹²¹. For example, the thought that '*a* is *F*' might carry certain entailments about the sorts of things that are and are not *F*, and the sorts of things that *a* is and is not; and the orthodox view of thought holds that such inferential relations are constitutive of that thought's content. So, on this view, it seems that attributing such a thought to a subject makes sense only if they can display a grasp of such inferential relations by, at the very least, being able to entertain the thoughts which are inferentially related to the original thought in such a way as to partially constitute that thought's content. And this is likely to involve, at the very least, the ability to think that other things are *f*, and that *A* has, or could have, other properties.

Considerations such as this motivate the proposal of Evans' 'generality constraint' (GC) as a hallmark of conceptual thought. GC demands that thoughts be closed under possible meaningful recombinations of their constituents – so if a subject can think 'Conor is Irish' and 'Ezio is Italian', they must also be able to think 'Ezio is Irish' and 'Conor is Italian' according to GC. Consequently, a thought's satisfying GC implies "it has a fine-grained intrapropositional structure that enables the subject to decompose and recombine its elements promiscuously and transfer them to other contexts, and to generalize and make quantificationally structured inferences that depend on such context-free decompositional

¹²¹ My characterisation of these issues is influenced by Toribio (ms): *Is Perception Compositional?*

structure.”¹²² Philosophers such as Peacocke (1992, ch.2), McDowell (1998) and Davies (1992) think that satisfying GC is a conceptually necessary condition of having thought, whilst Fodor (1987) holds that it is an empirical matter whether all thinkers satisfy GC, but that it is extremely unlikely that we will find a thinker that violates it¹²³.

All this begins to look problematic for enactive theories when we consider that the most obvious example of a domain where we find structured items that decompose and recombine in a general way, and can be inferentially related to each other, is natural language and its sentences, words and meanings. The most obvious way to account for an agent’s ability to think, reason and plan, on the orthodox conception, would be to appeal to their credentials as a competent user of a language. But a strength of enactivism was supposed to be its ability to account for the presence of consciousness in non-linguistic agents. Even if a defender of the orthodox conception thinks that abilities for thought can be explained without appeal to abilities to use a public language, they must still hold that the thinker has *language-like* abilities – they can decompose, recombine, and make inferences on the basis of, certain mental states just as if the contents of those mental states were sentences in a language. So even if the orthodox view stops short of requiring linguistic abilities of a thinker, they still have a cognitively demanding conception of the conditions that must be met in order to entertain genuine thoughts – a conception such that many intuitively conscious agents will fail to meet it.

The past few paragraphs have been motivating the idea that the capacity to entertain certain thoughts involves meeting GC, or some similar principles of de- and re-composability that can allow the requisite inferential relations between contents. Will any of this be relevant to the action-space theory, which emphasises only the availability of the action-oriented contents of perception to reasoning and planning? It seems so – until an enactive theory convinces us otherwise, the most natural way of construing reasoning and planning is as thought devoted to the topic of possible courses of action. And the role of the rational and inferential relations that motivate constraints such as GC seems particularly pertinent for thought in this domain. Reasoning and planning is naturally understood as involving counterfactual deliberation between a range of possible alternatives, and the relative evaluation of the likely consequences of each of the possible courses of action. In order to do

¹²² Hurley (2006) section 3.

¹²³ These references are from Beck (ms.)

this, it seems that an agent must grasp that their current situation affords a range of possible actions, the ways in which those actions relate to likely consequences, and the ways in which those consequences relate to the agent's current interests and projects. The same consequence can be a good one in the light of one project, but a bad one in light of another. And the same action could result in one sort of consequence in one environmental context, but a radically different one in another. The domain of reasoning and planning seems to provide a particularly stark example of a case where an agent must manifest sensitivity to complex sets of rational relations obtaining between both the constituents of both real and counterfactually entertained environmental situations, actions, outcomes and projects. And the orthodox view invites us to understand this ability in terms of the agent's being in cognitive states that it can systematically decompose, recombine and inferentially relate to each other, like the sentences of a natural language and their constituents.

Thus, according to a natural view of thinking, reasoning and planning, if an agent's grasp of the space of enabled actions that their perceptual environment currently affords is to be correctly described as figuring in her thought, her grasp must be such that she conforms to GC, or some similar principles of the possible de- and re-combinations of the constituents of her deliberations. This is problematic for the action-space account's claim that access to a space of enabled actions need not be a conceptual achievement. When setting out the action-space account, I claimed that neither the appreciation of currently enabled actions nor the integration of that appreciation with planning and reasoning required full-fledged, context-neutral conceptual abilities. A natural way of cashing this claim out would be an appeal to the possibility that the way in which a perceiver entertains the possibility of an action, the satisfaction of a goal, or the relations between those actions and goals and the perceiver's higher-level plans and projects, might fail to meet GC: an agent's perceptual sensitivity to a visually presented fruit might enable them to grasp that the fruit affords eating, whilst being unable to grasp that other objects to which they are perceptually sensitive do or do not afford eating, or that the satisfaction of other of their goals is or is not afforded by the fruit. An agent's grasp of the actions afforded by her environment was supposed to be a nonconceptual achievement since understanding it in terms of conceptual capacities has the counterintuitive consequence of ruling out many non-human agents and cognitively-impoorished human agents as seats of perceptual consciousness. But the material in this section motivates the worry that we lack a grasp of what thinking, reasoning and planning could be, if not capacities that implicate concept use. We can express the problem as an apparent dilemma for the action-space account – it either lacks an integration requirement, ruling in too many systems as seats of

consciousness, or includes such a requirement, but in doing so rules out too many agents as seats of consciousness. These conflicting pressures on the account have arisen because of a tension between two of its claims. On the one hand, we claimed that the perceiver's grasp of the ways in which she is poised to act is nonconceptual in virtue of her inability to flexibly decompose, recombine and make general inferences on the basis of, the constituents of the contents she grasps. But, on the other hand, the account requires that the same content function as input to thinking, reasoning and planning – all capacities that are naturally understood in terms abilities to decompose, recombine and make inferences on the basis of the constituents of the content grasped, which amount to the possession of conceptual capacities¹²⁴.

I think the enactivist should respond to this dilemma by agreeing that making sense of a content's availability for use in thought, reasoning and planning requires that content to be subject to certain principles of de- and re-composability, in a manner that respects the rational relations that obtain between possible contents. But they should deny that this amounts to requiring that conscious agents possess full-fledged, language-like, context-neutral conceptual abilities. We can make sense of planning, reasoning and thinking without requiring unrestricted recombining of the constituents of the contents involved in those processes, and thus appeal to a notion of practical reasoning that requires more than mere responsiveness to the environment, but less than full-fledged conceptual abilities. Our solution should make intelligible how an agent could be appropriately described as thinking, reasoning and planning without such abilities. And the account which enables it to do so should be constructed in a way consistent with the enactive focus of our view. The next sections develop a possible solution. A consequence of what follows is that many non-human and human thinkers will violate GC, held by some to be a conceptual impossibility. The next sections suggest a way in which this might be so.

5.2 Materials for a Solution: The Shared Circuits Model

¹²⁴ I'll be assuming this notion of conceptuality in what follows. It might be that there are reasons to view the less demanding form of reasoning, planning and thinking I'll defend in the subsequent sections as involving concept use, but I won't say anything about this issue. See Toribio (ms.) and Poston (2007) for arguments that states that cannot be decomposed into freely recombinable constituents can still be characterized as conceptual. For what it's worth, it seems to me that a sensible moral to draw from the account I develop in what follows is that we shouldn't think of conceptual thought as all or nothing – it comes in degrees, with limited flexible generality at one end, and full (GC meeting, linguistically enabled) generality at the other.

I'll argue that considering Hurley's (2004, 2008) 'Shared Circuits Model' (SCM) suggests how we can give the necessary response. In this section I'll outline the relevant features of the model, and in the next I'll say how it bears on our problem. Before that, some caveats. SCM is a deep and complex proposal, and I'll present a severely attenuated version of it, focusing only on the aspects relevant to my proposed solution. I'll thus be leaving out much of the empirical support for the model that Hurley draws from myriad sources. Multiple suggestions concerning how SCM ties in with work in cognitive science at various levels can be found in her original treatments. SCM is a thoroughly empirical and testable hypothesis, with an impressive level of detail. As such, it could easily turn out to be false in some or all of its aspects. The proposal I develop here relies only on fairly high-level features of the view, and so might still be a correct account of the way in which perceivers meet the integration requirement if many of the low-level details of SCM are incorrect. But I think my proposal has some value even if the aspects of SCM on which it relies turn out false, since it outlines *one way* (there may be others) in which an agent could meet the integration requirement in the absence of general conceptual abilities. Part of the challenge suggested in the previous section was to demonstrate that this was even an intelligible possibility. I claim that the account I offer does this, even if it turns out that agents do not in fact meet the integration requirement in this way. Now, on to the model.

SCM is proposed by Hurley as a model of subpersonal mechanisms that could enable personal-level capacities. A chief virtue of such a focus is the possibility that by framing proposals about how information is processed at the subpersonal level, we can learn about how things are at the personal level, just as knowledge of the micro-structure of an economy can give us new insight into, and the resources to better characterize, its higher-level structures. Hurley wished to use the model to provide enabling explanations of social cognitive phenomena such as action-understanding and mind-reading. We shan't be concerned with these features of SCM here – our project is to investigate the extent to which it can provide us with a model of counterfactual deliberation about possible actions that could inform our integration requirement.

To get a sense of SCM's focus, and the use we'll be making of it, note that much of its motivation is the provision of an account of imitation and imitative learning, and their significance. Imitation is a demanding cognitive capacity that humans are particularly

disposed to, and that we find only in cognitively sophisticated creatures; demanding because it requires an agent to perform a novel action to achieve a goal by copying another's goal-directed action. Thus it's distinct from the more common capacities for movement- or response-priming, where only the means to a goal is copied (as in the coordinated flocking behaviour of birds, or schooling behaviour of fish), and goal emulation, where only the goal, but not the means is copied (as in a case where an agent perceives an attractive goal of another's action, and tries to achieve that goal through their own means). Imitation requires the copying of both, and thus an implicit grasp of the means/end relation between them. Some such grasp of this relation is required since (for example) a monkey could use a tool to achieve a goal in the same way as she saw another monkey use the same tool for the same goal, but fail to be imitating – she may be copying only the action's end, and has coincidentally adopted the same means to pursue it, or copying only the means, but has found that it can be used to achieve an end that she desires, that happens to coincide with the end of the action originally copied. Imitation requires a grasp of the means/end structure of intentional action, and thus turns out to involve “something phylogenetically rare: the flexible interplay of copying ends and copying means: a given movement can be used for different ends and a given end pursued by various means.” (Hurley (2008), p.5).

Five different levels or layers of control mechanisms are described in SCM. At layer 1, a system has the capacity for adaptive control of behaviour to achieve a goal, using environmental feedback as guidance. For example, a thermostat has the goal of making the room a certain temperature. It achieves this goal by transducing an input signal – the room's temperature – and using a comparator circuit to compute the difference between the actual and desired inputs. We can say that this mode of control is an 'inverse model': one where the system employs a function that specifies a means, output or action on the basis of comparing actual and desired inputs. It thus 'maps goal to means' (Ibid. p.22), or specifies the means that will be used to achieve the goal. This sort of control is adaptive – different outputs are needed for the same goal in different circumstances, and the control mechanism allows the system to select outputs in light of the state of the environment. Importantly for our purposes (as will become apparent), Hurley hypothesizes that:

“Specific means/ends associations or instrumental mappings can be chained (output A is the means to controlled result B, while B in turn is the means to controlled result C, and so on) or organized into hierarchies. There are independently determined

evolutionary, developmental, and individual differences in the grain and complexity of the possible control sequences and hierarchies of different creatures.” (Ibid. p.20-1)

Thus, the control mechanisms a system uses for various of its actions might link up, or fail to link up, in various ways¹²⁵. An agent might have different subpersonal mechanisms controlling its grasping behaviour depending on whether it employs a power or precision grip, and another mechanism that controls its throwing behaviour (for example). The grasping and throwing mechanisms might link up – the power-grasping mechanism takes sensory inputs and, via use of an inverse model, converts them into motor outputs resulting in the grasping of the object. The endpoint of this control mechanism could serve as the starting point of another mechanism¹²⁶ – the throwing mechanism might take as input visual and/or proprioceptive information about an object held in a power-grip, and use an inverse model to form a motor plan detailing the movements required to throw the object in a particular way. But not all control circuits need link up in this way – there might be an agent who can throw objects held in power, but not in precision, grips. We could explain this by saying that the control circuit for power grips, but not for precision grips, is such that it can be linked to the circuit for throwing objects.

Though the feedback-modulated control at layer one is adaptive, it is also slow, since the feedback runs through the environment. Another drawback of such a minimal control system is the possibility of overshooting its target value – our thermostat produces an output signal that is appropriate to its goal of making the room a certain temperature, but the signal takes a while to have the requisite effect on the room’s temperature. In the meantime, the thermostat has received environmental feedback telling it that the room is still too cold, and it thus boosts its output signal, eventually making the room too hot. Layer two of SCM employs simulated prediction of effects for improved control. The system retains an ‘efference copy’ of its motor signal, and builds up an association between this signal and the feedback that is received as a result of it. As a result of acquiring these associations, a system’s motor outputs will evoke expected input signals. The mapping of output onto expected input in this way is a ‘forward model’. Our thermostat might improve its control process by employing a forward model in this way, predicting the consequences of its current output signal for future

¹²⁵ Making the same point, Hurley notes that “SCM doesn’t describe one all-inclusive structure, but has multiple instances for specific movements and results, at various points along different means/ends chains.” (Ibid. p.25)

¹²⁶ Consequently, we should understand the means/end distinction as relative rather than absolute – the same action can be an end for one control circuit, and a means for another.

temperatures, and modifying its output as a result. This can speed up and smooth out a system's control, since the system no longer needs to wait for environmental feedback to compute the next appropriate output, relying instead on simulations of the likely input¹²⁷. Such mechanisms could also be adaptive for a system by providing informational resources for distinguishing between actions on the part of the agent, and impingements on the agent by the world. Actions by the agent will be accompanied by predictions about likely effects, so the resulting input will be 'expected' by the control system. This will not be true of impingements by the world upon the agent¹²⁸.

At layer three, we suppose that the output/result associations built up at layer two can be activated bilaterally – as well as output signals evoking expected input signals, input signals can evoke likely motor responses. As a result, observed actions will be copied, if the evoked motor response is strong enough, or the system's capacity to inhibit it is weak enough¹²⁹. At this level, we have a control system that can enable "...simulation of means/ends associations from either direction: observed action retrodicts motor activation in the observer via mirroring of causes, which are associated with further results via simulative prediction of effects." (Ibid. p.24) Depending on the complexity of, and relations between the control mechanisms available to an agent, fairly complex forms of action and cognition could be enabled at this layer. For example, Hurley remarks:

"No doubt a monkey can move her hand to grasp a piece of sushi and move it to her mouth to eat it. But I can move my hand to operate chopsticks to pick up sushi to dip it in soy sauce and then move it to my mouth to eat it, in order to impress my boss; given associated simulative mirroring functions, I may start to resent you for eating the last piece of sushi as soon as you reach for your chopsticks." (p.26)

The idea here is that an agent might have separate control systems for object-grasping, grasped-object-eating, chopstick-wielding, sushi-grabbing, boss-impressing, and so forth. Such control circuits could be linked and articulated so that the input of visually-presented

¹²⁷ See Clark and Grush (1999, section 1) on skilled reaching behaviour for a nice example of this sort of control and its adaptive function.

¹²⁸ See Frith's (1992) explanation of thought and action-insertion in schizophrenia for a discussion of how deficits in such control processes could distort a subject's ability to distinguish between their actions and impingements from the world.

¹²⁹ For certain actions, this supposes that the system can solve the 'correspondence problem' of mapping the observed action of another on to possible actions of their own. See Hurley (2008) for a survey of suggestions about how this might be achieved.

chopsticks evokes a motor plan for sushi-grabbing, and the evoked sushi-grabbing is registered as something can be employed as a means to the social end of boss-impressing. When I see someone else reach for their chopsticks, such linkages between my control systems could enable me to perceive their action as a means to sushi-grabbing, and perhaps to perceive their predicted sushi-grabbing as a means to boss-impressing, and thus lead me to resent them. Hurley is contrasting these flexible relations between the ends of some control circuits and the means of others, and the sophistication of the sorts of control circuits we might have (e.g. for complex social functions such as boss-impressing) with the possible inflexibility and lack of sophistication we might find in the control apparatus of a monkey, and the corresponding simplicity of their intentional behaviour.

Layer four introduces the capacity to inhibit actual behavioural output, and monitor the fact that output is being so inhibited. This capacity is of particular importance for our purposes, since when combined with layer two this could enable offline instrumental deliberation. At layer two the use of simulations for improved control doesn't require the control system to distinguish actual and simulated feedback. If simulative predictions can be taken offline, and the fact that they are offline can be monitored by the agent, this could provide information about results of possible actions. As a result of such information, "Simulated results of alternative possible actions could be compared for the closest match to a target prior to actual action." (Ibid. p.28) thus enabling deliberation and choice among possible courses of action by the agent. For example, if my goal is to impress my boss, the respective expected inputs for output actions of sushi-dipping and racist-joke-telling could be compared for the best match with the goal state of my boss-impressing control circuit. Finally, at layer 5 inputs as well as outputs can be taken offline, enabling counterfactual simulation of inputs. At this stage, fairly elaborate counterfactual deliberation could be enabled. The counterfactual inputs of my boss smiling, laughing, or looking aghast could be entertained, and the different possible courses of action in light of each one could be compared in virtue of the apparatus at layer four.

A final thing to note about SCM is that Hurley intends it as an instance of a horizontally-modular organization of information-processing resources. That is, neither the ordering of layers nor the actions governed by these layers are to be understood in terms of a hierarchical, vertical organization. Rather, different sorts of actions will have control systems of different levels of sophistication, and different possibilities for links with other control systems. The fact that the control circuit for one of an agent's actions has the level of complexity we find at layer five does not entail that the control circuits for all her actions will. And the fact that

one of her control circuits is linked to many others does not entail anything about the way those other circuits might be linked to each other¹³⁰.

Now we have the necessary elements of SCM on the table, let's see how they might suggest the right way to construe the integration requirement.

5.3 The SCM Solution – Practical Reasoning Without Full Generality

The above consideration of SCM is relevant to our attempt to spell out the integration requirement since the way it enables rational relations between means and ends, and the possibility that means/end pairs can be decoupled and recombined suggests a less demanding analogue of the unrestricted de- and re-composability we find in language and concept use, and thus suggests the materials with which to state our integration requirement¹³¹. SCM hypothesises that such capacities could be enabled for an agent by a series of progressively more sophisticated subpersonal mechanisms for sensorimotor control. The prospect, then, is that SCM could show us how an agent's meeting the integration requirement (and thus enjoying perceptual consciousness, according to the enactive views) falls out of its credentials as a source of intelligent behaviour. Such an account of the integration requirement would thus meet one of our initial desiderata – stating the requirement in a way that meshes with the action-space account's emphasis on skilled, worldly action.

The key feature of SCM for our purposes is the way it suggests ends and means might be flexibly related for an agent, capable of de- and recomposing in different ways according to the levels of sophistication, articulation and connectivity of its control structures. Recall that the horizontally modular structure of SCM means that the level of complexity involved in a system's control circuits, and the level of connectivity obtaining between those circuits, need not be uniform throughout the system. Thus the structure of an agent's control system can impose limits on the way ends and means can be de- and recomposed for her. It is this point that suggests how the model can open up space between the full and free recombability that

¹³⁰ See Clark (2001, p.91) on Brooks' can-collecting robot for a good illustration of a horizontally modular architecture and its virtues.

¹³¹ This suggestion about the proto-conceptual nature of intelligent intentional action is found in various places in Hurley's work. However, to my knowledge, she never linked this suggestion up with her views on the enactive nature of experience. In 5.6 I indulge in some speculation about why this was so.

the orthodox conception of thought demands, and brute, unthinking responsiveness to the environment¹³². SCM suggests how thought devoted to reasoning and planning need not rely on full de- and recomposability of its constituents, with such flexibility and generality coming in degrees, according to the structure of the system's control apparatus. But before this proposal begins to look like a solution, we need to clarify a few points about how SCM, a proposal at the subpersonal, non-normative level, is supposed to relate to practical reasoning, a normatively constrained, personal-level phenomenon.

First, note that we're not principally interested in SCM as a story about subpersonal architecture – we're interested in the personal-level agency it enables, and how SCM suggests that the structure of that agency could be constrained by the structures of the agent's control systems. Thus, we're interested in SCM's role in enabling an agent to perceive something as a reason for acting a certain way, and as affording or failing to afford the satisfaction of certain of her goals. Does mapping our subpersonal sketch of SCM onto personal-level conclusions about the ways means and ends can be related for an agent involve an over-hasty projection from subpersonal to personal levels? We can give a negative answer to this question since SCM has been explicitly framed as an enabling hypothesis about personal-level action and perception. When we speak of the target value for a particular control circuit, we mean the value that corresponds to the achievement of a personal-level goal by the agent. The hypothesis is that such targets are tuned and re-tuned via multiple failed and successful interactions with the environment, so the subpersonal target value comes to correspond to the success of the personal-level action. SCM suggests how real and counterfactual perceptual stimulation can be used as input for the control circuit in charge of achieving a particular goal, and the results of this checked for their match with the target state of that circuit. In this way, SCM suggests an enabling hypothesis about how perceptual input can become meaningful for a subject by being processed in the light of its relevance for the pursuit of the subject's current goals. Just as the target values of a suitably environmentally calibrated circuit correspond to personal-level goals, the possible outputs of a control circuit (which can be taken offline at higher levels of SCM) come to correspond to possible intentional actions on the part of the agent, and the possible inputs to the circuit come to correspond to possible distal

¹³² This possibility seems to be anticipated in Hurley's suggestion that the recombinant structure of means and ends might be an evolutionary precursor of the recombinant structure of language (ibid, p.11), and suggested by the connections she alludes to between SCM and her work in *Making Sense of Animals* (e.g. Hurley (2006), section 6). The current suggestion is that the limited recombinability we find in the practical deliberation enabled by SCM might be a precursor of the full generality characteristic of thought according to the orthodox conception.

environmental events. Thus, SCM's role as a mechanism by which an agent can use environmental information to accomplish their intentional goals means that, so long as an agent's circuits are suitably tuned to the relevant actions, events and goals, we are entitled to move from claims about the subpersonal SCM to claims about the personal-level activity it enables¹³³.

These links between SCM and personal-level agency combine with the possibility that an agent's control circuits can be linked and articulated in various ways to suggest how the action-space account can appeal to something more than brute responsiveness to environmental affordances but less than fully general conceptual abilities when spelling out its integration requirement. Our SCM-inspired solution suggests that mere responsiveness to affordances fails to meet the integration requirement (and thus fails to suffice for conscious experience) since such responsiveness doesn't require a grasp of the means/end structure of action. An agent might perceive that a mug affords grasping, but without appreciating that grasping is an action that other objects afford, or that the mug affords actions other than grasping. They thus fail to distinguish the means and ends of their action, or the relation between those means and ends. The possibilities of links between control circuits (corresponding to the possibility of structured intentional action) and the counterfactual simulation of inputs and outputs (corresponding to means and ends) that we find at the higher level of SCM would enable the means and ends of an action to be distinguished, decomposed and recombined in new ways by an agent, in a way reminiscent but falling short of the full generality we find in conceptual thought.

Suppose it's granted that SCM makes it intelligible how ends and means could be flexibly decomposed and recombined by an agent in the absence of the full generality characteristic of thought according to the orthodox conception. An objector could still question whether such a

¹³³ I assume here that the right sort of calibration between the system and its environment and goals makes it appropriate to speak of parts of the system corresponding to environmental features and goals. Note firstly that this doesn't entail the claim that such correspondence *constitutes* a part of the system's carrying a certain content (though it might) – just that such correspondence contributes to enabling the creature underpinned by the system to be in a state with a certain content. Secondly, a creature's control systems having an SCM-like structure won't legitimate claims about its personal-level states unless that structure is sufficiently articulated and complex to support the decombination, recombination and appreciation of ends, means and their relations in a manner governed by principles of holism and normativity (see my discussion below of Hurley on how complex instances of SCM make it appropriate to characterise a creature's action in these terms).

story gives us anything worthy of being called even a primitive form of thought – surely, they might claim, all we have demonstrated is the possibility of relatively sophisticated control of action. And recall that part of the problem for the action-space account was that it needed an appeal to at least *something like* thought about a creature’s poise over an action-space to avoid an implausible liberalism in its attributions of conscious experience. For the action-space theorist to demonstrate that the variety of practical reasoning suggested by SCM enables them to frame an integration requirement in a way that avoids such liberalism, they need to say something about how the view of practical reasoning on offer relates to the orthodox conception of thought sketched in section one.

Worries about whether the view of practical reasoning here bears any relation to thought construed according to the orthodox conception would most likely be motivated by skepticism about whether our view can account for the rich normativity that characterises the relationships between thoughts, and between their constituents. Note firstly that the picture of control under consideration has various sources of normativity that stem from elements of the control circuits postulated by SCM possessing *correctness conditions*. The role of an inverse model (layer one of SCM) is to select a means appropriate to achieve a given goal, on the basis of comparing current and target inputs – insofar as selections of means can in fact be more or less appropriate to the achievement of a goal, they have correctness conditions. Similarly, the role of a forward model (layer two) is to predict likely sensory input on the basis of current output – insofar as these are predictions that can be more or less accurate, they have correctness conditions also. Finally, we suggested that the way that perceptions evoke associated actions according to SCM enables perception of the environment in terms of the sets of actions it affords. Such perception of affordances also has correctness conditions, insofar as the agent can be right or wrong about whether their current environment really does afford the action in question. However, as suggested above, the mere possibility of states with such correctness conditions does not suffice for the sort of intelligent agency to which the action-space account wants to appeal. Both sensitivity to affordances and implicit expectancies about the sensory consequences of action could be in place for an agent that merely responds in a fixed way to environmental stimuli. The key source of normativity for our model is the way in which an agent’s grasp of ends and means is such that the ends and means of a given action can decouple, and be flexibly recombined in ways that are sensitive to the ways in which they relate to the goals of the agent and the right ways to achieve those goals. This feature, as emphasized by Hurley (2006) makes it appropriate to characterize the behaviour of the agent

in terms of holism and normativity¹³⁴, and gets us a form of intelligence at least on a continuum with the orthodox conception of thought with which we begun this chapter. The higher levels of SCM enable a form of intelligent agency where:

“Perceptions and intentions combine to make certain actions reasonable and appropriate from the animal’s perspective. Means and ends can decouple: an intentional agent can try, err, and try again, can try various different means to achieve the same end. The holism of intentional agency provides a minimal, coarse kind of recombinant structure: an intentional agent has the ability to combine a given intention with different perceptions, given ends with different means.” (Ibid. Section 2)

Hurley’s idea is that the right sort of interplay between action and perception can make it appropriate to speak of reasons for action that apply from the animal’s perspective, and thus to begin to characterize the mental states of the animal that subserve that action in normative terms that anticipate the richer normativity and generality we find in the orthodox conception of thought. This is because in the sort of behaviour in question, relations between environmental stimuli and the agent’s responses are not invariant but (normatively) constrained by a context of means/end relationships, such that “A different intention or end will yield different actions given different perceptions about means/ends contingencies, and vice versa.” The agent thus views environmental inputs and behavioural outputs through the lens of their current goals, and their perceptions of the relations between means and ends. Thus, though an agent might be able to directly perceive that a visually-presented fruit affords eating, this perception might not result automatically in eating-behaviour – for a suitably sophisticated agent, the affordances of the fruit are seen in light of their mesh with the plans and projects it is currently engaged in¹³⁵. For example, if our agent is engaged in a

¹³⁴ I focus here on the holism and normativity that result from our SCM-inspired view since I assume that they are the features of intentional mental states which are most difficult and important for a theory of mental content to account for. See Haugeland (1990) for an excellent taxonomy of attempts to naturalise intentionality, anchored around the different ways of accommodating holism and normativity. See Fodor and Lepore (1992) for a more sceptical look at the links between holism and intentionality.

¹³⁵ As Hurley puts it, in a quote that sums up much of the material in this section, “Perceptual information may not lead to an invariant response, but explain action only in the context set by intentions and the constraints of at least primitive forms of practical rationality. Perceptions and intentions may combine to make certain actions reasonable and appropriate from the animal’s perspective, and mistakes are possible. As explained, the holism and normativity of

projectile-based food fight, the fact that the fruit is something that affords throwing will be more salient to its current activity than the fact that it affords eating. The different effects of throwing and eating the fruit can be compared for the best match with the target value of a control circuit in charge of the animal's food-fighting behaviour, and the action most appropriate to its current goals selected and performed on that basis.

Note firstly that this gets us a form of practical rationality – in light of her current projects, throwing the fruit rather than eating it really was the rational thing to do. And we might also think that the agent acts *for this reason* – if her control circuits have been suitably tuned by interactions with the environment, the fact that throwing fruit rather than eating it is the right thing to do when one is in a food fight explains and guides her action. We can say that these are reasons *for the agent* (that they are personal-level and normative, rather than subpersonal and causal) since, as noted above, we are interested in SCM chiefly as a hypothesis about the sorts of personal-level capacities that might be enabled by various forms of sensorimotor control – thus we are interested in the relative evaluation of counterfactual inputs to a control circuit in terms of this process's personal-level description of evaluating the ways possible means contribute to the satisfaction of a given end. And, as we have just seen, when the control apparatus of an agent is suitably complex, it can become appropriate to interpret the agent's actions in a way constrained by principles of holism and normativity – constraints which separate the personal level from the subpersonal. It thus becomes appropriate to speak of this process as giving the animal reasons for action, from her own perspective¹³⁶.

Secondly, note that the sort of process we have described relies on the agent being able to decompose and recombine means and ends, and be sensitive to the normative relations between them. The agent can see the fruit both as affording eating and as affording throwing, and the goal of winning the food fight as one that might be pursued via various different courses of action. In this case, our agent correctly decides to achieve this goal by throwing the fruit rather than eating it. But as we have been noting, the generality displayed by an agent in their practical reasoning can fall short of the full generality that comes with conceptual

intentional agency bring with them a kind of coarse recombinant structure, but this falls well short of enabling the context-free inferential promiscuity of conceptual abilities. An animal's various goals could nevertheless give him reasons to act in one way rather than another in particular circumstances--his own reasons, reasons from his own perspective." (Ibid, 6.9)

¹³⁶ Hurley (Ibid, section 6.9) notes that "The animal may not *conceptualize these as reasons*--but to require that would be to beg the question at issue," since the question at issue is whether there can be reasons for action without fully general conceptual capacities. I won't try to settle whether genuine reasons must be conceptualised as such here, but see Peacocke (2001) and Hurley (2001) on Brewer for arguments against this claim.

abilities. SCM can explain this in terms of the presence or absence of links between different control circuits¹³⁷. So, we have arrived at a suggestion of how an agent could display a form of practical reasoning that is normatively constrained, gives the agent personal-level reasons for action, and where the mental states involved are subject to at least some principles of compositionality. This shows us how our SCM-inspired model occupies a space between mere responsiveness to the environment and the orthodox conception of thought¹³⁸ – the space we needed to locate in order to give a plausible account of the integration requirement on behalf of the action-space view.

5.4 The Integration Requirement

So, what does the possibility of such an intermediate conception of practical reasoning suggest about how the action-space theorist should cash out their integration requirement? Recall that our previous gestures at such a requirement demanded that an agent's grasp of the actions currently afforded by her environment be such that it could be integrated with their ongoing practical reasoning, or that the agent have some kind of direct, non-inferential access to the space of currently-afforded actions. In the first section of this chapter, I suggested that a problem arose from these suggestions, since the natural construal of practical reasoning is in terms of the orthodox conception of thought, making the integration requirement too difficult to fulfill for creatures who lack general concepts, but are nonetheless intuitively conscious. Likewise, the most natural way to construe the requirement that an agent have 'access' to the space of actions they take to be currently afforded is in terms of Block's (1995) conception of access – the availability of a content for reasoning and report. These are capacities which also invite understanding in terms of the orthodox conception.

¹³⁷ I won't say anything here about the possibility of a SCM-inspired solution to another problem about normativity and content – how to tell whether an agent has successfully pursued one goal rather than failed to achieve another (slightly different) goal. This is the traditional rule-following problem that arises for any attempt to naturalise intentional content. Since it's not clear whether or how it can be solved at all, the fact that it arises here shouldn't count against our proposal (as also noted by Hurley (Ibid, section 6.4)).

¹³⁸ Making the same point, Hurley suggests that "Perhaps the metaphor of a space of reasons should be replaced with the metaphor of islands of reasons emerging from a sea of causes. For us human animals, language provides bridges that finally link these islands together." (Ibid, section 6.9)

But our consideration of SCM has suggested a less demanding way of meeting that integration requirement. We can appeal to a form of practical reasoning that does not rely on linguistically-mediated thought, but consists in being able to entertain the possibilities of various different actions and grasp how these actions would contribute to the satisfaction of the agent's current goals. So when the action-space account claims that an agent's grasp of the possibilities for action afforded by their environment must be such as to allow practical reasoning about those possibilities, this should be understood as requiring that the agent be able to grasp that the afforded action is a means to one among a selection of ends, and the relationship between their performance of that action and the satisfaction of their current goals and projects. Thus, whilst DF's spared perceptual sensitivity to the orientation of the slot in the posting task might prime her to behave in ways that can be brought out via prompting, she fails to be conscious of the slot's orientation since she lacks an automatic grasp of the ways she is currently poised to act relate to her current goals and projects (of engaging in a slot-posting task). Standard perceivers possess such a grasp, and are thus able to use the information about the slot's orientation to achieve their current goal. Sophisticated agents can practically reason about whether performing the posting task, or engaging in some other currently-afforded action yields the best chance of satisfying their *overall* goals and projects.

To take a more complex example than the simple slot-posting task, suppose someone is playing a game on the Nintendo Wii, where they must orient their wrist in such a way, and move their hand appropriately in space, to match the orientation and location of an avatar of themselves that appears on screen to the orientation and location of an avatar-shaped slots in a series of bubbles that float towards the bottom of the screen. The object of the game is to pop the bubbles by posting your avatar through each of their slots before they hit the bottom¹³⁹. This game is relatively easy for a standard perceiver, since perceptual information about the orientation and position of the slots in each bubble is automatically put in touch with their current goal of preventing the bubbles from hitting the bottom of the screen. According to the action-space account, this explains their being conscious of the positions and orientations of the bubbles. More cognitively sophisticated agents could engage in practical reasoning about which bubble to pop first, based on comparison of the anticipated consequences of popping different bubbles. On the other hand, according to our interpretation of her condition in chapter two, DF would be unable to play this game successfully. Though her brain still processes information about the orientations of the slots in each of the various bubbles before

¹³⁹ Interested parties can find this game on the Wii-play disc for the Nintendo Wii. Thanks to Mog Stapleton for bringing this game to my attention.

her, this information cannot be automatically used to guide her performance in the game in the way that it can for a standard perceiver. If she were prompted to begin popping-behaviour when the game began, her success in prompted slot-popping tasks makes it likely that she could succeed in popping some bubbles. But such responsiveness to the bubble-popping affordances of her environment does not suffice for consciousness according to the action-space view, since (I conjecture) DF's responses to the ways in which the falling bubbles afforded a avatar-popping cannot be put in touch with her overall goal of popping the bubbles before they reach the bottom of the screen

It's perhaps useful to compare our formulation of the integration requirement with Hurley's (1998) notion of 'nongognitive access' to information. A creature has such nongognitive access when processed information is automatically available for use in their forming an intention whose content is provided by that information (Ibid. p.149), where 'automatically' is understood in such a way as to rule out reliance on prompting or forced-choice conditions¹⁴⁰. For example, someone with prosopagnosia, a deficit in their ability to recognise faces, "does not have the normal ability to form an intention such as: if the face is familiar, say 'hello'...And she is unable to spontaneously form correct intentions to act on the basis of information about whether the face is familiar." (Ibid. p.152) A theory that makes nongognitive access to content a necessary condition for a conscious state's having that content says that the access she lacks is necessary for consciousness of facial similarity. But in doing so, we have not tied intentional access, or the ability to make use of processed information, to conceptual abilities:

"Information can be used explicitly to meet a creature's immediate needs, even though it is used in a context-bound way and the creature lacks the ability to use general concepts of itself, its states or the objects or properties of which it is conscious in a variety of contexts detached from those needs." (Ibid. p.153)

Our account of the integration requirement plausibly coincides with a requirement of nongognitive access to contents for consciousness. The action-space account claims that the contents of conscious experience are action-oriented, and that for a content to become conscious, it must be apt for use in the subject's current practical reasoning. We have tried in this chapter

¹⁴⁰ Hurley also includes the condition that the agent be able to act successfully upon this intention. This seems too demanding - for example, it appears to rule out the possibility that a temporarily-paralysed agent could have conscious experience whose content and character is explained by his ability to form intentions upon which he could not act.

to cash out this metaphor of ‘aptness’ in SCM-inspired terms of an agent’s abilities to recognise an action as a particular means to various possible ends, and to grasp the relevance of relationships between means and ends for the satisfaction of their current goals. On a plausible view of what it is for an agent to have an intention, such abilities are criterial – brute responsiveness to the environment is not intentional behaviour, since true intentional behaviour requires responsiveness that manifests sensitivity to the way the means required to achieve the goal of the behaviour vary according to context, and the ways in which the same means might be used to pursue different goals in different contexts. And accounting for the possibility of such responsiveness without requiring full conceptual abilities was the purpose of our appeal to SCM, and the basis for our integration requirement. Consequently, our integration requirement says that for an action-oriented content to be part of the content of a conscious experience, that content must be available to guide the agent’s intentional behaviour, where ‘intentional behaviour’ requires sensitivity to a range of possible relationships between means and ends, as brought out by our discussion of SCM and the relation of the abilities it enables to the orthodox conception of thought.

5.5 Thinking about Colour

How does our statement of the integration requirement relate to our account of colour experience in chapter four? Recall that there we gave an account of the content of colour experience in terms of an agent’s implicit knowledge of the range of discriminations currently afforded for them by their perceptual exposure to a coloured object. Our integration requirement states that for information to become part of the content of a conscious experience, that information must be available to inform the content of an intention on the part of the agent. We might think there are a couple of sources of tension between our account of colour experience and our statement of the integration requirement. Firstly, it seems that the sorts of capacities in terms of which we explained colour experiences, such as the capacity to discriminate between different shades of colour, or to locate a colour in a relationally-defined colour space, are not actions we seem to engage in very often – certainly not as often as we experience colour, a pervasive feature of our visual consciousness. Secondly, our account of the integration requirement has appealed to the idea of control circuits for the sorts of afforded actions that constitute the content of conscious experience according to the action-space theory. This appeal allowed us to emphasise the way the means and ends of a given circuit might decouple, and serve as possible inputs or

outputs to different control circuits – and it was argued that possibilities such as these let us see how reasoning about intentional action can be a normatively constrained and personal-level phenomenon without requiring linguistic or conceptual abilities, and thus meet the desiderata for our integration requirement. But is it plausible to think of our discriminations of colours as structured and controlled in this way? If so, what are the possible means and ends that might decouple and recombine in colour perception? Addressing these questions will help clarify both our account of colour experience and the integration requirement.

In response to the first potential source of tension we should note that our requirement has not demanded that our ability to make certain discriminations must *in fact* contribute to the content of a current intention in order to become the content of a conscious experience, merely that our grasp of the fact that such abilities have been enabled be *apt* to contribute to our current intentions. Thus, the fact that we perceive colours when we are not engaged in matching or discrimination tasks need not be in tension with the integration requirement. Our discussion of the action-space account in general in chapter three made it plausible that abilities being apt for such integration was importantly linked to the content of a subject's experience – recall our analysis of the colour-blindsight patient (chapter 4.2). Such a patient can reliably make fine-grained colour discriminations between colour samples presented in their blind-field, but only when externally prompted, and sincerely report lacking any experience of the colours in question when doing so. The action-space account claims that this is because they fail to meet our integration requirement – their discriminatory abilities are not such as to be automatically put in touch with their intentional behaviour (in this case, performing a discrimination task). This is not so for standard perceivers. When absent-mindedly casting my eye over the range of coloured objects before me, the colours of those objects form part of the content of my visual experience, since my perceptual sensitivity to them is such that I could at any time decide to (for example) intentionally categorise a seen object according to its colour, or pick the background which provides the best match for that colour¹⁴¹.

¹⁴¹ In addition to our general arguments from chapters 3 and 4, demonstration of inattentional/change blindness for colour (see e.g. Chabris and Simons (1999)) seems to support our claims about conscious colour vision here. Large changes in the colour of a visually-presented object can go unnoticed when a subject is attending closely to other features of the object or scene. The action-space account might explain this in terms of the irrelevance of the object's colour to the agent's current projects. Whilst unnoticed changes in colour might prime subsequent responses or performance, they are not consciously experienced if the discriminatory abilities enabled by them were not at that time apt to contribute to the content of an intention formed by the subject, perhaps due to the way the subject's attention was focussed at the time. However, the correct interpretation of change and inattentional blindness results is

The second potential source of tension is more complex. Broadly construed, we can see it as a worry about the apparent disanalogy between the types of actions we appealed to in our account of colour vision and those appealed to when setting up our account of the integration requirement. One sort of disanalogy is that the actions appealed to in colour are discriminatory and apparently automatic, whereas those appealed to in our discussion of SCM are bodily and intentional. In response, note firstly that an object's colour will often make a subtle background contribution to our intentional behaviour. The way an object looks with respect to its colour helps us to distinguish it from other surrounding objects, perceive its shape via the patterns of light and shade displayed on its surface¹⁴² and even to perceive what kind of object it might be. According to the action-space account, such abilities and dispositions aren't straightforwardly *results* of our perceiving a particular colour – rather they help constitute what it is to perceive that colour. In this way, abilities to act upon objects based upon their boundaries, shape and even their semantic properties are bound up with experience of colour, and such abilities look more like the intentional bodily abilities to which we appealed when setting out our integration requirement, undermining the apparent disanalogy.

This helps us see how to respond to the point that the actions to which we appealed in our account of colour experience don't seem to have the kind of means/end structure that was crucial to the actions we appealed to in our discussion of the integration requirement. Whilst discriminating one colour from another is not naturally understood as an intentional action with a goal, and a means by which that goal is achieved, our account of colour perception makes dispositions towards intentional actions that *do* have such structure, such as sorting objects based on their colour, or acting on an objects based on the shape or texture they appear to have, constitutive of our colour experience. More importantly, on our account of colour experience, there are analogies between the structure of intentional action and the structure of the abilities in terms of which we explain colour experience which it is illuminating to note.

The two-level structure of our account of colour perception, according to which perceived patterns of colour-looks suggest the colour a surface is likely to have, and the colour a surface is perceived to have suggests the patterns of colour looks it is likely to display, is reminiscent

controversial (see e.g. Block (2007)), and I won't attempt to argue for this interpretation in further detail here.

¹⁴² Recall that our account of colour made grasp of the significance of such patterns partly constitutive of perceiving an object's colour (Chapter 4).

of aspects of our discussion of SCM and the view of intentional action it suggests. The way perceiving a surface as a particular colour entails implicit expectations about how it would look in different lighting-conditions and contexts looks like the implicit predictions about likely consequences of action that forward models run. The way in which perception of the current array of colour sensations from the surface entails a particular perception of its colour looks like the inverse models that take a system from registration of current input to predictions about the likely upshot of that input for possible actions. And the way in which understanding colour is a matter of appreciating the relationships between possible sensory inputs and possible behavioural outputs thus looks like the way in which SCM suggests that we come to understand the actions of ourselves and others, and the world's possibilities for action.

We saw that we could move from talk of relationships about inputs and outputs in SCM to talk about relationships between means and ends, and it seems that an analogous move is available here. In the same way that one movement will be a means to a different end in different contexts, colour sensations are a means to different colour-categorisations in different contexts – the colour-look of the squares in the checkerboard illusion (4.3) could correctly result in its being judged to be either light or dark grey, depending on the context in which it appears. One pattern of sensations can thus indicate different surface colours, and one perceived surface colour can also be indicated by different colour looks – just as one end can be achieved by a variety of different means, depending on context. Understanding the relationships between colour-looks and colour properties is key to perceiving colour, just as understanding the relationships between possible means and ends is key to being conscious, on our account.

Finally, just as the ways in which the means and ends of possible actions relate to each other is governed by constraints of holism and normativity, so is the way in which colour-looks, surface colours and the discriminatory abilities which account for our perception of each, are related. A surface's displaying a particular array of colour-looks makes it appropriate for a perceiver to categorize it in one way, but not in others, and a perceiver's categorizing a surface as a particular colour makes it appropriate to expect it to display certain colour-looks as viewing conditions change, but not others. And on our account, a creature's ability to perceive one colour is holistically related to its abilities to perceive all others, in a way dictated by the geometry of its colour space (recall the sudoku comparison from the previous chapter), in the same way as a creature's ability to perform one intentional action is related to its abilities

to perform actions with the same means but a different end, and vice-versa. But just as the generality implicated in intentional action falls short of the full generality demanded by GC, and characteristic of thought according to the orthodox conception, the way the discriminatory abilities constitutive of colour-perception relate to each other is also subject to certain restrictions. Beck's (ms.) discussion of cognitive capacities that violate GC but are naturally described as involving thought suggests a story about the nature of these restrictions. Beck's starting point is work which shows that pigeons can reliably make discriminations between two quantities only if the difference between those quantities is sufficiently large. Pigeons seem to be able to think thoughts such as "40ish button-pecks differ from 50ish button-pecks", and "38ish button-pecks differ from 47ish", but not "37ish button-pecks differ from 40ish", or "47ish button-pecks differ from 50ish". It turns out that, for pigeons,

"...the ability to discriminate two quantities is a function of their ratio. For example, just as the pigeons can discriminate 43 from 50 but not 47 from 50, they can also discriminate 86 from 100, but not 94 from 100. To put it formally, their numerical discriminations obey *Weber's Law*, $\Delta^I/I = k$, where I is the value of the magnitude (in this case number), k is a constant, and Δ^I is the minimal change in magnitude required for discrimination."

(Ibid. p.5-6)

Thus, the constituents of these thoughts are de- and re-composable only in a limited manner that violates GC. These results are of general interest, since it turns out discriminations in many different modalities and domains, such as duration, rate, size, weight, loudness, brightness, pressure, taste, smell, pain, length, area and distance, have Weber constants for humans as well as pigeons. Thus some sorts of thoughts about these quantities will violate the generality constraint, since their constituents will be decomposable and recombable only subject to the limits imposed by Weber's law.

I don't pretend that the remarks in this section are conclusive – these ideas all require working out in further detail. But I think they do enough to allay the worry that there is a problematic tension between the action-space account of colour perception and our integration requirement. The key points of this section were firstly that the actions to which our account of colour-perception appeals do not have to be executed for their enabling to form part of the content of an experience, and that the way the colour-constituting abilities figure in reasoning and planning will most often be as a background constraint on the sorts of intentional actions

available to an agent. Secondly, I argued that there are more parallels between the structure of intentional agency and the structure of the abilities at the heart of our colour experience than we might expect, and that the disanalogy is thus not as starkly problematic as it might first appear.

6. Consciousness, Self-Consciousness and an Enactive Pact?

On our account, perceptual consciousness and a certain kind of practical reasoning ability are tightly linked. I want to end this chapter by suggesting that this feature of the account might tell us something about the relationship between consciousness and self-consciousness, and the relationship of the action-space theory to Hurley's sensorimotor view. There are two distinctive features of conscious states that we might think a theory of consciousness must explain: firstly that such states have phenomenal properties - there is *something-it's-like* to be in them; secondly, that such states are always had from a particular perspective - to explain what a conscious state is like is to explain what it is like *for someone or something*. So far in this thesis I have focused on only the first feature, but it seems plausible that each of these are essential features of conscious states - if there is nothing it is like to be in a state, then that state is not conscious, and if a state is not such as to imply a particular perspective on the world, that state is not conscious. Our account of the integration requirement in this chapter, and our discussion of the explanatory gap in the next, suggest how these features of experience are related. Recall that it was required of our account of practical reasoning that it yield reasons for action that apply from the perspective of the agent, rather than merely causal reasons for a system's behaviour. In section three, I argued that certain sophisticated forms of control involve sensitivity to the relations between possible means, possible ends and the current goals of the agent, and that these relations are both holistic and normative in character - holistic because the same means can be used to achieve different ends in different contexts, or given different goals of the agent, and normative because of the ways in which a given means can succeed or fail to be appropriate for achieving a given end¹⁴³. This meant that certain forms of control can provide reasons for action from the agent's perspective in the absence of fully general thought - thus, our account of practical reasoning suggests how personal-level states that implicate a particular perspective on the world can come into

¹⁴³ Recall from section three that there were other normative relationships between means, ends and contexts besides this one.

existence by virtue of certain forms of sensorimotor control. If such a story is on track, our action-space theory can account for the essentially perspectival nature of consciousness.

It might be thought that this is the most important feature of consciousness that a theory must explain - since it is not clear whether the possibility of states with phenomenal properties that are not had from a particular perspective is one we can make sense of, we might be skeptical of a theory that explains only the phenomenal properties of experiences without offering an explanation of their perspectival nature, since it seems to admit this possibility. However, there are also reasons for thinking that the explanation of the phenomenal properties of experience should take priority – we might agree with a story about the conditions required for a system to have a unified perspective on the world, but question whether the presence of such a perspective entails the presence of consciousness, and is thus question whether we have been given a story about ‘self-consciousness’, rather than one about non-conscious selfhood. For example, Bermudez’s account of self-consciousness (1998, 2001) identifies various sources of self-specifying information, such as the essential reference of somatic proprioceptive sensations to one’s body, and the fact that the visual field is bounded, manipulable, and flows towards a central point when motion occurs, and uses them to ground an account of how it is possible to entertain thoughts with self-referential content in the absence of linguistic or conceptual capacities. But it might be objected that whilst such appeals can account for the presence of states carrying information about how things are within a certain bounded system, this does not explain why such information should issue in any awareness, or consciousness¹⁴⁴. It seems possible that a simple robot might process information about the states of its parts, or be able to recognise patterns in the flow of input from a video camera as it moves around its environment, and do so in the absence of any conscious experience. Even the presence of the more sophisticated forms of control to which we have been appealing, involving compositionality of means and ends seems compatible with the absence of conscious experience, for similar reasons. We might agree that such control structures involve a normativity and holism that makes it appropriate to say that the system in question can act inappropriately or appropriately in light of its current perceptual and motivational states –

¹⁴⁴ I don’t claim here that such an objection applies to Bermudez’ account as a whole – just to an account that emphasised *only* such informational properties of a system.

but this seems only to suggest that we should view the system as an intentional one, leaving the question of whether or not it is also conscious open¹⁴⁵.

It seems, then, that there are motivations for viewing each property of conscious states as primary. It might be thought that this leads to another apparent dilemma – a story about the phenomenal properties of experience will not be satisfactory unless it explains the perspectival nature of experience, and a story about the perspectival nature of experience will not be satisfactory unless it explains the phenomenal properties of that experience. We cannot claim to have a theory that explains either one unless we first have a theory that explains the other. But such a problem would only arise for theories that attempt to account for phenomenal properties whilst bracketing the obligation to give an account of the perspectival nature of experience, or vice versa. The action-space account, however, is not such a theory. This chapter has argued that taking an enactive focus allows an explanation of how a perspective on the world can arise from gradually more sophisticated forms of sensorimotor control. The following chapter argues that the same focus also allows us to see how to address the explanatory gap. Thus, the conditions that the action-space account states are necessary for phenomenology are also necessary for a perspective on the world, and vice-versa. The action-space account can therefore explain why the perspective of an agent with suitable control apparatus should involve phenomenology. And the fact that the necessary conditions the action-space account gives for conscious experience involve an integration requirement, and meeting the integration requirement involves (as we have seen) a particular perspective on the world, means that the action-space account will not admit the counter-intuitive possibility of isolated and unowned phenomenal states.

It seems plausible that Hurley never (to my knowledge) made an explicit link in her work between her sensorimotor account of consciousness and her work on practical rationality for reasons suggested by the above discussion. For Hurley, questions about the perspectival unity of consciousness and its phenomenal character were tackled together. Her work on practical rationality overlaps with her work on the unity of consciousness, which is in turn bound up with her views on the interdependence of perception and action, and the role of that interdependence in an account of consciousness. She thus never made an explicit link between her work on practical rationality and her work on consciousness since, I suggest, she did not

¹⁴⁵ These are abbreviated versions of the explanatory gap worries I considered in more detail in chapter one – the reasons for thinking that such exercises in conceivability are important were covered there.

view these as properly distinct topics. So it would not have occurred to her to use the model of practical reasoning suggested by SCM to provide an integration requirement for her sensorimotor view of consciousness since in developing such a model, she already took herself to be addressing the question of how a creature can come to have a conscious perspective upon the world. This might cause us to wonder whether the action-space view and Hurley's view as I have depicted them here are really distinct. For each depicts an agent's consciousness as arising from a certain kind of practical knowledge of the possibilities pertaining to action in its environment, and cashes this practical knowledge out in a way inspired by Hurley's work on practical rationality. In 1.5 we contrasted action-space and sensorimotor views by noting that action-space views emphasise knowledge of the possible consequences of perception for action, whilst sensorimotor views emphasise knowledge of the possible consequences of action for perception. But the model of practical reasoning set out in this chapter seems to cast doubt on whether such a contrast can be sharply drawn. Recall that according to our model of practical reasoning the distinction between ends and means is not absolute – the same movement can be an end in the context of one goal and a means in the context of another. We might likewise think that, on our model, the distinction between knowledge of the sensory consequences of action, and knowledge of the practical consequences of sensation is not absolute. For according to it, perceptual inputs are understood in terms of their significance for possible actions (*this* sort of perceptual state affords *that* sort of behaviour), and possible actions are understood in terms of their likely effects on perceptual inputs (*that* sort of behaviour affords *this* sort of resultant state) and then evaluated according to the match of the predicted inputs with the agent's goals. Thus we can employ neither a one-sided emphasis on the consequences of action for perception (as Noë does), nor the consequences of perception for action (as the action-space account prior to our introduction of the integration requirement does), since our model suggests that the grasp of perceptual consequences is in terms of their significance for action, and the grasp of practical consequences is in terms of their significance for future perception. This suggests that the contents of perceptions and intentions are co-enabled, co-dependent and co-defined, as Hurley argued in *Consciousness in Action*.

So does our action-space view merge into Hurley's two-level interdependence view with the introduction of the integration requirement? I am not convinced that it does – it still seems open for a proponent of the action-space account to agree that the contents of perceptions and intentions are co-enabled and thus co-dependent, but differ in holding that appeal to an implicit grasp of a space of available actions plays a privileged role in the explanation of

experience. I suggested in this section that a theory of consciousness must account both for the perspectival nature and the phenomenal character of conscious experiences. Hurley's sensorimotor view suggests an account of the first feature, but is largely silent on the second. Nothing in her view suggests how we might construct either a conceptual bridge between phenomenal properties of experience and sensorimotor dynamics, or why such a conceptual bridge is not required in an explanation of phenomenal consciousness. We saw in chapter one that a successful explanation of consciousness must meet one of these obligations. Hurley restricts her explanatory aims with regard to consciousness to addressing the "comparative" rather than the "absolute" explanatory gap¹⁴⁶ – the question of why one experience differs from another, rather than why there is any experience at all. Whilst I think it is likely to be true that a specification of the features in virtue of which one experience differs from another will include reference to sensorimotor dynamics, I don't think that bracketing questions about why experience occurs at all is a legitimate way of explaining experience. For any two differing experiences, there will be numerous candidates for the property or properties that explain that difference. My visual experience of a cube will differ from my visual experience of a sphere in terms of its sensorimotor dynamics, but also in terms of the range of actions I implicitly take to be enabled in each case, the levels of activation of certain brain areas that help to enable that experience, the physical states of certain neurons that make up those brain areas, and the physical states of the particles composing my body as a whole and the objects in my environment. How can we decide which of these levels of explanation should be privileged in our account of the difference between the two experiences? It seems to me that the only way to adjudicate between these competing candidate levels of explanation is to provide reasons why one of them should be more relevant to the character of experience than others, and that until we have a story about why *any* non-phenomenal level of description is relevant to the character of experience, such reasons will not be forthcoming. The next chapter attempts to provide those reasons on behalf of the action-space account.

If the position argued for there is successful, we have reason to think that whilst both directions of explanation between action and perception are relevant to constructing an account of perceptual and intentional content and practical reasoning, the action space direction of explanation takes precedent when explaining the phenomenal character of experience. In keeping with the conclusions of the previous chapters, we would thus endorse an action-space account of the content and character of experience, whilst retaining something like Hurley's

¹⁴⁶ E.g. Hurley and Noë (2003), Hurley (forthcoming)

two-level interdependence view as an account of how the relevant poise over a space of actions is enabled.

Conclusion

I began this chapter by arguing that the action-space account's appeal to an agent's ability to think, reason and plan about the affordances of their environment disclosed by their perceptual sensitivity creates a dilemma. On the one hand, such an appeal is necessary to avoid ruling in very primitive information-processing systems as seats of perceptual consciousness. On the other, since thinking, reasoning and planning are standardly understood as cognitively demanding abilities, requiring something akin to linguistic, or general conceptual, abilities, the appeal looks to rule out all but the most cognitively sophisticated systems from the ranks of conscious perceivers. The action-space account, we worried, was too liberal without such a requirement, and too conservative with it. However, via a consideration of Hurley's shared circuits model, we saw that there could be forms of practical reasoning that go beyond brute responsiveness to an environment, but stop short of requiring the linguistic and conceptual abilities thought to be largely confined to mature human beings. I argued that when an agent consciously perceives that some action is currently afforded by her environment, she grasps that the end whose satisfaction is afforded might, in other contexts, also be satisfied by different means, and that the means that the environment currently solicits from her might, in other contexts, be used to achieve different ends. Such an interplay of means, ends and contexts can bring with it a holism and normativity that fall short of the richer normative relations that characterise linguistic and fully general conceptual thought, but can still support cognition that is appropriately described as practical reasoning. An agent with a suitably rich and flexible repertoire of actions can grasp the different ways in which currently afforded actions could contribute to the satisfaction of her various goals, and settle on a course of action on the basis of a comparison of the predicted consequences of those ways. The action-space account claims that, to the extent to which she can do this, she enjoys conscious experience.

The final three sections of this chapter explored some consequences of this account of the integration requirement. In 5.4 I suggested that, upon a plausible conception of intentional action, our integration requirement requires with Hurley's (1998) suggestion of 'intentional

access' to contents as a necessary condition on a content's being conscious. In 5.5 I showed how our account of the integration requirement meshes with chapter 4's account of colour-perception. Lastly and, in my view, most importantly, I argued in 5.6 that developing the action-space account in this way opens up the possibility that an account of self-consciousness could fall out of our account of consciousness. Tying consciousness to a capacity to practically reason about intelligent intentional action, as suggested in this chapter, ties consciousness to a particular, normatively constrained and coherent perspective on the world. In doing so, we account for what many have thought to be an essential property of conscious states – that they are always the conscious states of someone, or something, had from a particular perspective. But I also suggested that questions about the perspectival nature and the phenomenal character of conscious experience must be tackled together. For without an account of the relation between knowing poise over an action-space and the phenomenal character of experience, it might justly be asked why our story so far constitutes a theory of consciousness at all. Additionally, without such an account, there is no story about self-consciousness in the offing from the action-space approach – our account in this chapter might hint at how particular perspectives on the world come into existence, but it would give us no reason to believe that there should be anything it is like to occupy such a perspective. Much hinges, then, on whether the action-space account as I have sketched it can provide a way of responding to the challenge of the explanatory gap outlined in chapter 1. Fortunately, I think that it can, and showing this is the task of our next chapter.

Chapter 6: The Action-Space Theory and the Explanatory Gap

It is time to address a question that has been looming throughout our discussion of the action-space theory – why does the action-space account deserve to be called a theory of *conscious experience* rather than, say, of perceptual content, or of cognitive capacities that usually coincide with conscious experience? The theory claims that conscious experience consists in a subject's grasping the space of enabled actions disclosed to her by her perceptual sensitivity to the environment in such a way that she can practically reason and form intentions on the basis of that sensitivity (as we saw in the previous chapter). But how does this help us with the problems with which we began, in chapter 1? Could we not imagine a perceiver's possessing such a grasp whilst lacking conscious experience, or two perceivers manifesting identical grasps of identical spaces of enabled actions, but nonetheless differing experientially? An upshot of the considerations of chapter one was that if a putative account of conscious experience cannot say something about these questions, then we have been given at best a story about plausible functional or intentional correlates of consciousness, which falls short of a true *explanation*. However, I wish to argue here that the action-space account can make progress on these questions – it can tell us why the situations they invite us to conceive are in fact unintelligible, whilst explaining why they do not strike us as such, and why we are tempted to worry about such questions in the first place. I'm going to do this by surveying three popular ways of addressing the explanatory gap – appeals to cognitive access, representationalism, and the phenomenal concept strategy – and arguing that the action-space account lets us see what the methods have in common, and how they may be fruitfully combined. Once we see the common basis of the three views, we can see how the criticisms that can be levelled at each should be resisted.

6.1 The Gap Revisited

This section briefly reviews material covered in greater detail in chapter 1. In 1.1 we saw that the alleged explanatory gap between the phenomenal – the domain of facts about what it is like for subjects to have particular experiences – and the physical stems from our inability to specify a conceptual connection between physical and phenomenal facts about persons, and our consequent inability to move from facts about the physical way things are with a person to

the ways things are for them phenomenologically¹⁴⁷. The zombie, inverted spectrum and knowledge arguments we considered there were all intended to bring out the intuition that we lack a grasp of how the physical and the phenomenal might be conceptually linked. For example, the knowledge argument works by inviting us to imagine a subject, Mary, with the maximal amount of non-phenomenal information about the mind and its workings, as well as unlimited abilities to draw inferences and conclusions from this stock of information – if anyone were able to move from physical to phenomenal facts, it would be Mary. However, we are additionally invited to suppose that Mary has spent her life incarcerated in a monochromatic environment. The popular intuition is that, despite her complete physicalistically-acceptable knowledge and deductive powers, Mary will nonetheless lack knowledge of certain facts about what it is like to have certain experiences – in particular, visual experiences of colour. This suggests that we can see neither a conceptual connection between the phenomenal and physical realms, nor the possibility of such a connection – for if we could, we would conclude that Mary in her monochrome room would possess information about what it was like to see colour, due to her ability to move from the physical information she possesses to phenomenological conclusions.

We also saw that demonstrating the conceptual mesh between explanandum and explanans that seems lacking in the case of the physical and the phenomenal looks like a necessary condition on a successful reductive explanation. Employing Chalmers' taxonomy of attempts to naturalise phenomenal consciousness, we saw in 1.2 that each position appears unsatisfactory. Type-A materialists hold that there is in fact a conceptual mesh between phenomenal and the physical realms, but do so in the face of strong countervailing intuitions from the zombie, inverted spectrum and knowledge arguments – if reflection on our experience showed that physical properties can be conceptually linked to phenomenal properties, we would not have the intuitions that absent or inverted qualia were possible, or that Mary would lack knowledge of phenomenal properties in her black-and-white room. It seems, then, that the type-A materialist must employ a revisionist conception of experience, at odds with that suggested by our introspective evidence, if phenomenal properties are to be conceptually connected to structural and functional ones.

Type-B materialists hold that the physical necessitates the phenomenal, but that the relevant necessities are not ones which may be deduced *a priori*. Below, I criticise an

¹⁴⁷ Recall that my use of 'physical' facts abbreviates the domain of all facts that a naturalistically acceptable explanation might appeal to, including functional and intentional facts.

interpretation of the action-space theory as a type-B materialism that reviews the reasons for our dissatisfaction with such a position in some detail, but the essential problem was that type-B materialists might be able to tell us which physical processes give rise to consciousness, but could not tell us why or how they do so. This is because the necessities to which the type-B materialists appeal are *a posteriori* – if they were *a priori* the nature of our concepts would make it clear why and how the requisite physical conditions could not occur without certain phenomenal conditions. It seemed plausible to grant that understanding why and how the domains in an explanation relate to each other is a desirable feature of a reductive explanation, and one that the type-B materialist cannot deliver.

Similar considerations rule out a dualist reductive explanation of consciousness, since dualism is typically motivated by the claim that there can be no conceptual connection between the physical and the phenomenal. Lastly, type-C materialists hold either that our cognitive limitations mean that we cannot grasp the conceptual connection between the phenomenal and the physical, or that some kind of conceptual revolution is required before such a connection will be available.

I argued in 1.3 that there is room for an acceptable reductive explanation of consciousness in between type-A and B materialisms. We can agree with the intuition that it seems natural to construe our introspective evidence in such a way as to suggest that there can be no conceptual link between physical and phenomenal properties, but argue that an alternative construal of that evidence is available – one according to which all properties of our experience follow conceptually from structural and functional properties. Such a deflationist position avoids the standard criticisms of type-A materialism as eliminativist, whilst retaining its commitment to the existence of a conceptual mesh between the phenomenal and the physical. It also agrees with the type-B materialist that there is no reductive explanation with an *a priori* component of the phenomenal, standardly construed, in terms of the physical in the offing, but disagrees that we must conclude from this that only an *a posteriori* link between the phenomenal and the physical can be demonstrated. Such a deflationism thus incorporates the intuitions behind type-A and type-B materialism, without the features which made each of them unacceptable as candidate forms of explanation. For a deflationist strategy to be plausible, then, it must accommodate the introspective evidence that leads to the putative explanatory gap, but plausibly redescribe that evidence in a way that renders it compatible with physicalism. The redescription must be such as to explain why the anti-physicalist construal of the evidence is natural, but mistaken, and must be empirically plausible – otherwise the

deflationist has shown only how the explanatory gap *could* be closed (perhaps a worthwhile achievement in itself), but not how it can be closed for *our* experience. We also saw in 1.4 that a deflationist strategy should say something about the properties of conscious experience that are standardly thought to make it a problematic candidate for reductive explanation – ineffability, privacy, direct availability and intrinsicity. Crucially, the deflationist need only explain why our introspective evidence makes it natural to construe experience as possessing those properties, and may do so without granting that our experience *in fact* possesses those properties.

Chapter 1 argued that the above deflationist position is the form a reductive explanation of phenomenal consciousness should take. In this chapter, I will argue that the action-space account suggests an explanation of consciousness that meets all the requirements of the deflationist position we have set out. Before moving on to our consideration of how the action-space account can address these obligations, let's illustrate them by looking at three ways in which the action-space account might be thought to address the explanatory gap that strike me as lacking. Firstly, we might interpret the action-space account as an attempt to operationalise consciousness. Considerations such as our discussion of DVS in chapter 2, or the neat ways in which the action-space account can handle the problem cases of chapters 3 and 4 might prompt us to think that knowing poise over a space of enabled actions is at least a close approximation to the functional correlate of consciousness. Given this, perhaps we should see the action-space account's contribution to a solution to the explanatory gap as specifying a functional profile with which conscious experience can be plausibly identified. But this will not meet the constraints on a solution we have laid out. Such a strategy is an instance of the type-B materialism criticised in 1.2. For a reminder of why this position fails to address the explanatory gap, recall Jackson's (2003) argument that naturalising explanations without an *a priori* component are indistinguishable from eliminativism about the target explanandum. If science informs us that the temperature of a gas turns out to be its mean molecular kinetic energy then, without some conceptual link between temperature and mean molecular kinetic energy, it is indeterminate whether we should interpret this as a new discovery about temperature, or as the discovery that there was no such thing as the temperature of gases, only certain functions of their kinetic energy. Similarly, if a scientist were to inform us that conscious experience in fact consisted in quantum-level dipole oscillations in thousands of neural microtubules resulting in the collapse of a superpositional state to a post-reduction outcome state of consciousness, it is indeterminate whether we should understand this as new

information about consciousness, or as the discovery that there was no such thing as consciousness as we conceived it, only various quantum phenomena. This is because there is no clear conceptual path between such quantum phenomena and the experiences we wished to explain. Were there such a path, we would be able to see that the relevant quantum effects entail the phenomenal explanans, and the explanatory gap would have been closed. So, according to our standards here, finding a physical or functional correlate of phenomenal states does not constitute an *explanation* of those states unless something further can be said about the conceptual connections between the two levels¹⁴⁸.

Secondly, recall the strategy employed by sensorimotor theorists that we considered and rejected in 5.6 of distinguishing between ‘absolute’ and ‘comparative’ versions of the explanatory gap. The idea was that an appeal to sensorimotor dynamics might not afford an answer to questions of why an agent has some experience rather than none, but can afford an answer to the question of why they have one type of experience rather than another. Perhaps the action-space account could limit its explanatory aspirations and take such a tack. However, I argued that the problem with this proposal is that there are indefinitely many levels of description at which similarities and differences between the physical correlates of experiences can be expressed. A visual experience of a cube will differ from a visual experience of a sphere in terms of its sensorimotor dynamics, but also in terms of the spaces of actions enabled by the subject in each case, in terms of the neural activity correlated with each experience, and perhaps in terms of the computational states subserved by the collapse of the superpositional states of thousands of neural microtubules that correlate with each experience. How should we decide which of these ways in which the experiences differ should be privileged in our explanation of their different phenomenal characters? A natural suggestion seems to be that we should privilege any option that seems to have some conceptual connection to phenomenal consciousness. If all the options are alike in lacking such a connection, and in being merely correlated with the phenomenal level we wish to explain, then all options appear on an equal footing with respect to their credentials as candidate explanations of the differences between the two experiences. If, however, one of the options can tell us something about its conceptual connection to phenomenal consciousness, this gives us both a reason to favour it over the others, and addresses the ‘absolute’ explanatory gap – physical conditions that have a conceptual connection to phenomenal conditions would show

¹⁴⁸ In 1.2 we also considered the argument that explanations that appeal to identities without a conceptual bridge seem to entail a conceptual dualism that rules out the possibility of successful reductive explanation.

us why those conditions obtaining for a subject entails that the subject enjoys some experience rather than none. Thus, we can conclude, if a proposal is to address the comparative explanatory gap, it must first address the absolute gap.

Finally, it might be claimed that the action-space account as developed thus far affords a direct response to the conceivability arguments used to motivate the putative gap between the phenomenal and the physical. Proponents of the explanatory gap will claim that we can conceive of the enabling of the sorts of skills the action-space theorist appeals to in the absence of experience, and that this conceivability reflects the lack of a conceptual path from enabled skills to phenomenal consciousness. We might respond by appealing to the richness of the enabled skills, the fact that our grasp of such enabled skills need not entail our ability to conceptualise or report upon the precise nature and extent of those skills, and the requirement that a perceiver's grasp of the space of enabled actions be such that it can inform her ongoing practical reasoning. When we appreciate all these points fully, is it perhaps not obvious that we can conceive of a suitably rich body of skills being practically grasped in the appropriate way whilst phenomenal consciousness is absent¹⁴⁹. But such a response does not adequately address the challenge posed by the explanatory gap as we have sketched it. As we saw in chapter 1, the conceivability claims typically adduced by proponents of the explanatory gap are merely a means to bring out the apparent lack of conceptual path between structural and functional descriptions and phenomenal consciousness. Appreciating this shows us that merely undermining the conceivability claim need not banish explanatory gap worries – to meet these worries directly, we need to show that we have undermined it *by* providing a conceptual path from our structural/functional account to the appropriate characterisation of our experience. The response under consideration does not do this. Consider again the proposal that conscious experience was explained by quantum-level dipole oscillations in thousands of neural microtubules resulting in the collapse of a superpositional state to a post-reduction outcome state of consciousness. It's not clear that it can be objected of such a proposal that we can conceive in appropriate detail of all these conditions obtaining, but of conscious experience nonetheless being absent, or different. But this is not because the proposal has closed the explanatory gap – rather, it has put a set of conditions in place that are sufficiently complex

¹⁴⁹ It might be objected by some that they are perfectly capable of conceiving of the conditions set out by the action-space account being satisfied in every detail, and phenomenal consciousness being entailed. It seems to me that our necessary and sufficient conditions are multifarious and complex enough that I am unable to hold them all in mind, in full and at once, to see what does and does not self-evidently follow from them. It is possible that things might seem different to someone with a more expansive mind, however.

that it is no longer clear what is and is not consistent with their obtaining. What makes the above quantum explanation of consciousness unsatisfactory is not straightforwardly a matter of what can and cannot be conceived of consistently with the conditions it specifies, but how unfamiliar talk of microtubules and dipole oscillations is supposed to link up with the phenomenal explananda each of us is so intimately acquainted with¹⁵⁰. Again, this brings out the fact that an adequate explanation of conscious experience must build a conceptual bridge, or at least say something about how their proposed explanation relates to such a requirement. Only if this is done can the action-space undermine the true worries behind the conceivability arguments – if we have a conceptual bridge between the physical and the phenomenal, it will be evident to us which features of the physical conditions we have set down for conscious experience rule out their conceivably obtaining in the absence of consciousness, and how they do so.

As noted, I believe that the action-space account can meet the demanding constraints to a solution of the explanatory gap we have set out here. We can begin to see how this is so by considering its relations to some existing promising proposals for tackling the gap.

6.2 Three Ways to Bridge the Gap

In this section, I'll consider three promising proposals for tackling the explanatory gap. For each, however, I will raise two questions. Each sort of account aims to specify a set of naturalistically-acceptable conditions that will suffice for phenomenal consciousness. The first sort of question for each account is whether we have reason to believe that these conditions obtain in our own case. If not, then even if the account in question is correct, we will be able to see only how the explanatory gap could be closed for some hypothetical class of conscious perceivers, but not for us. The second sort of question for each account concerns how the conditions they propose relate to the deflationist framework for which I argued in chapter 1,

¹⁵⁰ Although what is wrong with this quantum proposal can be cast in terms of conceivability – the explanatory gap challenge could be met if phenomenal states self-evidently followed from the quantum computational states to which our proposal appeals, and their absence was thus strictly inconceivable if such computational states were present. This is simply a further reminder of why thought experiments were supposed to test for the presence of a conceptual bridge in the first place. The problem with using our action-space or quantum proposals to block conceivability arguments in the way under discussion is that this rules out the presence of a conceptual bridge – this strategy claims that the conditions we have put in place are sufficiently opaque that it is unclear what does and does not self-evidently follow from them.

and our requirement that a successful explanation of consciousness must build a conceptual bridge. I will argue that the action-space account allows us to see the connections between the proposals, and in doing so, suggests answers to the questions I raise for them. Our conclusion will be that the action-space account suggests a solution to the explanatory gap that incorporates the best elements of the three proposals we will consider, and that we have good empirical reasons to believe that the conditions it sets out for phenomenal consciousness are met by conscious human perceivers.

6.2.1 Access and the Zombie-free Zone

Andy Clark (2000) has argued that certain patterns of access to perceptual discriminations can entail phenomenal consciousness – that is, the presence of phenomenal consciousness might be entailed by information being available for a subject’s capacities for report, recall, and reasoning in a certain way. To motivate this claim, he asks the reader to imagine a creature (let’s call her Nicky) who can reliably make a range of perceptual discriminations, such as identifying or distinguishing objects based on their olfactory, visual, and tactile properties. As we noted above, there seems to be no entailment between the mere presence of such abilities and phenomenal consciousness. But now suppose that Nicky can report on her discriminations, and respond to our interrogations about the way she makes those discriminations. Asked about how she has just successfully discriminated between two stimuli, Clark suggests her response must fall into one of two categories:

“[She] must say *either*:

a) I have no access to the act by means of which I detect the differences. The answer just comes to me. I perceive nothing when I make my judgements – I simply find myself saying that there are two objects, one red and one yellow, and so on.

Or:

b) I have access not just to the products of my sensory activity, but also to certain aspects of that activity itself. For example, I am non-inferentially aware that I am using a visual rather than a tactile modality. I am aware that I see, rather than hear or feel, the difference.” (Ibid. p.30)

The presence of a kind of limited, non-inferential access to the nature of her perceptual discriminations would force Nicky to respond in the second way. Given the right sort of access to facts about how she makes her discriminations, she would automatically know when she

has made a discrimination by sight rather than by touch (since her access to those facts is non-inferential), but be able to say no more about the ways in which these modes of sensing differ (since her access to those facts is limited).

Furthermore, it seems that such patterns of access obtaining will result in Nicky claiming when we interrogate her that *feels different* to make discriminations by sight and touch. Since she has only a limited access to the features in virtue of which those modes of discrimination differ that need not entail an ability to conceptualise or report on those features, it would be natural for her to say that there is 'something it's like' to make discriminations by sight, as opposed to by touch. If Nicky had *no* access to the features in virtue of which the ways of sensing differed, then she will not claim a difference, in terms of what-it's-like or otherwise, between perceiving (say) the size of an object by sight and by touch. If she had *complete* access to the features in virtue of which the ways of sensing differed, then it is plausible that differences between sensory modalities would strike her only as differences between the content and extent of the information gleaned in perception, and we need not suppose that such a difference in content need *feel* like anything to her¹⁵¹. But if Nicky has the limited, non-inferential access suggested above, there will be a salient difference between discriminations she makes by sight and those she makes by touch – a difference that she can report and reflect upon, but can give us no further information about, save that those ways of sensing differ for her.

It seems, then, that the presence of some limited non-inferential access to features of information-processing has forced Nicky into a 'zombie-free zone', where what-it-is-like for her differs according to whether she is sensing by sight, touch or smell. Importantly, it does not look like the conditions we have used to demarcate this zombie-free zone presuppose conscious experience – we began by ascribing Nicky the ability to make various sorts of perceptual discriminations, added the ability to make reports about those discriminations, and gave Nicky's capacities for reasoning, reflection and report a limited degree of access to the states underlying those discriminations. Taken individually, we can conceive of each of these capacities obtaining in the absence of conscious experience. But when they come

¹⁵¹ Clark (2000) does not mention the possibility that the perceiver he considers has complete access to the features in terms of which its modes of sensing differ, which would require a third option to be included in the possible responses it could give in response to our interrogation. There are thus two ways that a creature's access to its sensory operations could fail to result in phenomenal consciousness – if it has too much or too little. Clark has emphasised this in recent presentations of this material.

packaged together for one perceiver, Clark's argument pumps the intuition that things will seem a certain way for that perceiver as she makes her discriminations – seeing, hearing and touching will each feel different to her, though she will not be able to say much to us or herself about how or why this is.

However, as noted above, I want to raise two sorts of issues for this proposal, and each of those that follow. The first issue for Clark's strategy is that we need to say something substantial about the sorts of patterns of access to which we are appealing – what exactly is it that is being accessed, and what does a perceiver's standing in such an access relation consist in? Once we have answers to these questions, we need to consider whether we have good reasons to believe that the relevant patterns of access obtain in our own case¹⁵². Without such reasons we would be in the unusual position of seeing how the explanatory gap could be bridged, but being unable to apply this method to our own conscious experience – possibly an advance on our current position, but surely not what we were after.

The second issue is that we might wonder exactly how Clark's story squares with the challenge posed by the explanatory gap as I have set it out above. The task was to specify a conceptual path from structural and functional facts (here facts about patterns of access) to facts about experience. Should we see Clark's argument as an attempt to provide such a conceptual entailment, or as an intuition pump designed to assuage the need for such an entailment? We can bring this issue into focus by considering an obvious objection to Clark's proposal, which he attributes to Dave Chalmers:

'Our argument, it will be said, 'establishes at most that a certain kind of access implies a certain tendency to judge and report a difference (the sort of tendency my zombie twin might have) rather than implying a real phenomenal difference.'" (Ibid. p.32)

It will then be objected that it was not our tendencies to judge and report that we wished to explain, but the conscious experiences thought to lie behind those tendencies. Clark's response is to note that his proposal also suggests how such tendencies can be justified or mistaken, since a subject might be mistaken about the access she has to the differences between her modes of

¹⁵² This question, as well as the analogous questions raised for each of the positions considered in this section, should not be taken as criticism of any of these proposals, since each aims only at establishing a general framework from which to address the explanatory gap, not fleshing out the details of such a framework, as my question demands.

sensing, or make a false judgement on the basis of that access. But, Chalmers will surely respond, the truthmakers of our judgements and reports about experiences are not nonphenomenal differences between modes of sensing, but the conscious sensory experiences themselves. Should we interpret Clark's account as building the kind of conceptual bridge between those nonphenomenal differences and experiences that proponents of the explanatory gap require, or as showing us that the demand for such a conceptual bridge is misguided? I think we are best able to see what Clark should say here by considering the relation of his account to our next two strategies.

6.2.2 Representation and Instantiation

A common response to the explanatory gap consists in claiming that demonstrating the existence of the gap relies on some kind of mistaken construal of the properties of experience we wish to explain. For example, Jackson (2003) argues that we should not conclude that physicalism about phenomenal consciousness is false on the basis of his knowledge argument (1982), and that representationalism lets us see how to avoid doing so. Start by assuming the truth of representationalism about experience – experience is exhaustively representational, meaning that any change in the character of an experience is a change in the way that the experience represents things to be¹⁵³. If representationalism is granted, this allows us to see that there are a couple of ways in which we can describe what happens to Mary upon her release from the black-and-white room. The key point to remember is that to represent something as being the case does not entail that there is an instantiated property that corresponds to the way things are represented as being¹⁵⁴. It is a mistake to think that saying “unicorns do not exist” commits us to some entity of which non-existence is predicated. The sentence should be understood as representing things as being a certain way, rather than picking out an object and predicating a property of it. Likewise, it is widely agreed that saying “I have an experience of red” does not entail that there need be some existent red thing to which I stand in an experiential relation. Assuming otherwise was the mistake of the

¹⁵³ See e.g. Tye (2000), Lycan (1996) for defences of representationalism. Beginning a defence of a representationalist solution to the explanatory gap with an assumption of the truth of representationalism appears question-begging, but, as we shall see, the argument aims only to prove the conditional that *if* representationalism is true, *then* we can see how the worries that motivate the explanatory gap rest on a mistake.

¹⁵⁴ See Harman (1990) for a longer defence of this claim. Harman's position there can be interpreted as an instance of the strategy under discussion in this subsection.

sense-data theorist's response to the argument from illusion¹⁵⁵. According to the sense-data theorist, our experience having the properties it does is explained by our acquaintance with sense-data that are appropriately related to those properties. Representationalism suggests an alternative picture – the properties which account for the character of our experience are representational properties. My having an experience of red is a matter of my being in a state that represents things as being a certain way. But my representing things in this way need not entail that I stand in a relation to some existent object with the represented property. To think this is so, according to representationalism, is to confuse a representational or intentional property (or intensional property, in Jackson's terminology) with an instantiated one.

How does this bear on what we might say about Mary upon her release from her room? Jackson's point is that drawing an anti-physicalist conclusion from the fact that Mary has a new experience when she leaves the room relies on a certain conception of experience. The anti-physicalist suggests that when Mary sees her first red object, she learns about a new property of experience (phenomenal redness) that the physical information she assimilated in her black-and-white room did not tell her about. But viewing matters from a representationalist perspective allows us to question this conception of experience. The above remarks showed us that we need not think of Mary's experience of red as involving her standing in a relation to some instantiated experiential property that physicalism does not tell us about. Rather, we can understand her as being in a new kind of representational state, one that her previous black-and-white environment rendered off-limits. The intuitive line of resistance to this idea is that merely saying that Mary represents things in a new way leaves out the fact that she learns something of the form "red things look like *this*". But whilst it is true that this is something Mary might say upon entering her new representational state, moving from that fact to the falsity of physicalism requires interpreting the above '*this*' as picking out some instantiated property of experience that is new to Mary, precisely the characterisation that the representationalist rejects¹⁵⁶. Once the possibility of this view of

¹⁵⁵ For more of this sort on the relation between sense-data theory and the mind-body problem see Crane (2000), whose position there also has affinities with the general strategy considered in this subsection.

¹⁵⁶ Of course, this view of matters is not incompatible with representationalism if the new property picked out is understood to be Mary's property of being in a state with a certain representational content. But this can't be the sort of acquaintance with a new property the advocate of the knowledge argument has in mind if the argument is to work against physicalism.

matters has been opened up, Jackson suggests we attempt to explain phenomenal consciousness by giving an account of the distinctive features of the representations which constitute it.

I take Rorty (1979, ch.1.2) to be suggesting a similar position in his discussion of the phenomenal and the physical: "Why do we think of the phenomenal as immaterial? We do so because, as Ryle put it, we insist on thinking of having a pain in ocular metaphors – as having a funny sort of particular before the eye of the mind" (Ibid. p.31). Rorty thinks that the phenomenal seems incongruent with the physical only when shift our focus from phenomenal states that are essentially properties of people to phenomenal particulars:

"The neo-dualist who identifies a pain with how it feels to be in pain is hypostasizing a property – painfulness – into a certain sort of particular, a particular of that special sort whose *ess* is *percipi* and whose reality is exhausted in our initial acquaintance with it." (Ibid, p.30)

But once we are viewing pains as phenomenal particulars with which people can be acquainted, rather than as a certain kind of state of a person, physicalism is doomed. For the essence of any putative phenomenal particular is exhausted by the way it appears to subjects who are acquainted with it, and thus any attempt to reductively explain such particulars in terms of factors other than how things appear to subjects will fail to capture the essence of that particular, to explain what is ultimately distinctive of it. Luckily for physicalism, as our discussion of Jackson suggested, this view of matters is not mandatory – we must take care to think of pains as properties of people, not as particulars, if physicalism is to be defended. Jackson and the representationalists propose that pains, and all other experiential properties, are properties of how the subject represents things to be.

Again, two classes of questions are raised by this proposal. Firstly, adopting this representationalist framework for a solution obliges us to provide a further story about the character of the representations to which we are appealing – do we have reason to think that representational states with whatever features are necessary to house an explanation of the character of phenomenal consciousness obtain in human perceivers? Secondly, one might also think that Jackson's argument does not completely dispel the residual intuition that merely representing things in a certain way need not entail phenomenal consciousness – we can grasp the details of his proposal and still, it seems, worry that the right sorts of representational

states might be present whilst phenomenal consciousness is absent. Fleshing out our account of the relevant properties of the representations constitutive of phenomenal consciousness might help with this – but can we say anything more generally about why a representational state should entail phenomenal consciousness? As with Clark’s proposal, I think we need to address this second issue by considering how Jackson’s proposal relates to the challenge of the explanatory gap as we have sketched it – is representationalism supposed to provide a conceptual bridge between the phenomenal and the physical, or to show that such a bridge is not required? And, as with Clark’s proposal, I will argue that we can see how the representationalist should respond by considering the relations between their position and the others considered in this section.

6.2.3 The Phenomenal Concept Strategy

The last general strategy for tackling the explanatory gap that I want to consider is the phenomenal concept strategy (PCS). PCS acknowledges that there is, on one interpretation, a principled gap between our conscious experience and any structural or functional account we might give of that experience. However, this gap need not stand in the way of physicalism, nor need it rule out the prospect of our having a good explanation of phenomenal consciousness. This is because, PCS argues, the presence of the gap is explained by our distinctive way of thinking about phenomenal states – by appeal to the special character of our *phenomenal concepts*, and the way in which that character differs from that of the structural and functional concepts in terms of which we wish to frame an explanation.

There are various types of phenomenal concept strategy (PCS) on the market¹⁵⁷, and we shan’t consider any of them in detail here. We are interested here in the general strategy and how it combines with the two proposals above and the action-space theory, rather than which detailed version of the strategy should be pursued. What unites the various instances of PCS is the claim that the concepts we use to think about our experiences – our phenomenal concepts – are *conceptually isolated* from our structural and physical concepts. Different types of PCS give different accounts of this conceptual isolation – for example, following Loar (1990) and Tye (2003) we might think that phenomenal concepts are conceptually isolated from structural and functional concepts since they are recognitional concepts, applicable only in the presence of the phenomenal state they pick out. In addition to this recognitional aspect, such concepts must

¹⁵⁷ E.g. Loar (1990), Hill (1997), Perry (2001), Aydede and Guzeldere (2005), and Papineau (2006).

also lack logical or inferential ties to our other concepts to be conceptually isolated in the requisite way¹⁵⁸. We might have a partially structural/functional concept of red that is inferentially related to objects, surfaces, light and so forth, but our phenomenal concept of red applies to our experience, and lacks such relations (Carruthers (2000), Carruthers and Veillet (2007)).

If it is granted that conscious perceivers possess phenomenal concepts that are isolated in the requisite way, then we can see how to resist the moves from the conceivability of zombies, the knowledge argument, and the intuitive presence of an explanatory gap between the phenomenal and the physical, to anti-physicalist conclusions. The conceivability of zombies is explained since the structural and functional concepts we use to give our description of the zombie are inferentially isolated from our phenomenal concepts. When considering whether a zombie is conceivable, we use a phenomenal concept to think about an experience, then consider whether the conditions for applying that concept must be fixed by the structural/functional description of the zombie we have put in place. Given that phenomenal concepts are (for example) conceptually isolated recognitional concepts, it will always seem conceivable that a zombie could lack the experience a phenomenal concept picks out, regardless of how our structural and functional concepts say things are, and will never seem that our structural/functional description allows us to infer the phenomenal way things are.

The same points apply to Mary's inability to move from structural and functional information to phenomenal knowledge of the experience of red. The structural and functional information she is given is conceptually isolated from the phenomenal information she is trying to deduce. Consequently, the relevant phenomenal information cannot be arrived at via inferential moves from structural and functional facts. For Mary to think 'red is (phenomenologically) like *this*' in the relevant way, she would have to deploy the phenomenal concept of red, and our account stipulates that this concept is only applicable in the presence of the state which it picks out, rather than on the basis of any amount of structural or functional information. Since Mary's environment has been designed so that state picked out by the phenomenal concept of red will not occur in her, she cannot token that phenomenal concept, and thus cannot know what it is like to experience red.

¹⁵⁸ This characterisation is taken from Carruthers and Veillet (2007). It seems plausible that the isolation requirement should only apply to connections between phenomenal and structural/functional concepts, since it seems plausible that phenomenal concepts (e.g. of colour experiences) might stand in relations of relative phenomenological sameness and difference to each other.

Finally, the intuitive appearance of an explanatory gap is explained by the structural and functional vocabulary on one side of our candidate explanation being conceptually isolated from the phenomenal vocabulary on the other side.

PCS can thus accommodate all our anti-physicalist intuitions about consciousness. But if we also assume that our possession of conceptually isolated concepts can be given a reductive explanation consistent with physicalism, we can see that these intuitions can be explained without conceding physicalism's falsity.

Again, at least two sorts of questions arise for PCS. The first concerns how the appeal to a class of concepts that are conceptually isolated from our structural, functional and other naturalistically-acceptable concepts is to be cashed out – how does such conceptual isolation come about, and do we have reason to believe that the phenomenal concepts of human perceivers may be isolated in such a way? The second question concerns what sort of solution to the problem of the explanatory gap we should see PCS as affording – should it be understood as a means of closing the gap by providing an explanation which meets the requirements of section two, above, or only as a means of avoiding anti-physicalist conclusions whilst accepting the existence of the gap? The discussion above suggests that we should see it as the latter – if our way of framing the explanandum when thinking about the explanatory gap involves deploying a phenomenal concept, then we cannot hope for an explanation of this in structural and functional terms, since these classes of concepts are isolated from each other¹⁵⁹. But might there be a way to retain PCS's attractive explanation of our intuitions regarding consciousness whilst defending the kind of deflationist solution to the explanatory gap argued for above?

6.3 Building our Bridge

¹⁵⁹ Carruthers and Veillet (2007) note that whilst the truth of PCS would mean that we could not infer the falsity of physicalism from (for example) the conceivability of zombies, PCS is neutral as to whether conscious experience is subserved by physical properties alone. An anti-physicalist could agree that introspecting a phenomenal state involves deploying a conceptually-isolated recognitional concept, but still hold that those concepts pick out intrinsic and non-intentional properties of experience that turn out to be incompatible with physicalism. Again, I will suggest that we can see how to use PCS to argue for physicalism when we view it in relation to the other proposals of this section, and in the context of the action-space approach.

In this section, I argue that the three proposals considered above can be combined in a way that resists the criticisms to each. Before doing so, however, I want to suggest that considering the relation of the action-space account to each of the above proposals suggests how they should each respond to the first sort of question I raised for them – how should the conditions they identify as sufficient for phenomenal consciousness be fleshed out, and do we have reason to think that such conditions are met by conscious perceivers?

The relation of Clark's (2000) proposal to the action-space account is fairly direct. In the original article, he leaves the question of just how to characterise the limited, non-inferential access to which he appeals open, but at one stage floats the following suggestion:

“Perhaps, then, what we have access to when we have access to the modality involved in the act of detection is the specific battery of skills that we could have deployed. Insofar as the sets of skills differ according to the modality involved [...] access to the sets of skills which could have been deployed would constitute direct non-inferential access to the modality in use...” (Ibid, p.35)

Our consideration of the empirical (and philosophical) support for the action-space account in previous chapters suggests that this is exactly right – we have been arguing that we should construe conscious visual experiences as an agent's perceptual sensitivity poising her over a space of actions currently afforded by her environment in such a way as to allow a form of practical reasoning about the afforded possibilities. Different perceptual experiences are characterised by a different range of enabled abilities and skills, and we saw in chapter 5 how an agent's grasp of the actions currently afforded by her environment could be direct and without conscious inference, apt to be integrated with ongoing practical reasoning, but nonetheless unavailable for cognitively demanding capacities such as verbal report and conceptualisation. The action-space account, then, suggests how an agent could have the kind of limited but direct access to features in virtue of which sensory episodes differ alluded to by Clark, and our consideration of the DVS data, the plausible diagnoses the action-space account affords of various pathologies of consciousness, and our reflections on colour vision all provide reasons to believe that such patterns of access obtain in conscious human perceivers, and are implicated in visual experience.

We can see how the support we have marshalled for the action-space account gels with a representationalist account of experience if we gloss the action-space account (as I sometimes have) as in part a thesis about the content of perceptual experience. We can understand a perceiver's being poised to act in certain ways by her perceptual sensitivity to the environment as a matter of her occupying a representational state, with the content of the state given in terms of the abilities which it empowers – the way that state represents the affordances of the perceiver's environment. Our account of the integration requirement in chapter 5, and the above discussion of Clark's proposal and the action-space account, suggests that this is the sort of content to which we have direct, non-inferential, but limited, access to in experience. Jackson (2003) and the representationalists are thus glossed by the action-space theory as claiming that we are prone to conflate this access to the action-oriented contents of our perceptual experience with acquaintance with irreducible phenomenal properties. To the extent to which we have seen support for the action-space conception of visual experience, we have also seen support for the claim that visual experience involves states with this kind of intentional content.

This also suggests how the action-space account might provide a well-motivated way of fleshing out a phenomenal concept strategy. As we saw in chapter 5, the action-space account demands that an agent's grasp of the actions afforded for her by her environment be such that those affordances can inform her practical deliberation – such integrability is necessary for her perceptual sensitivity to affordances to issue in conscious experience. We also argued for the possibility that an agent could grasp affordances in this way without being able to report, conceptualise, or make general inferences on the basis of them. Thus, the type of grasp of affordances that the action-space account requires of a conscious perceiver comes apart from the way affordances are grasped when brought under linguistic concepts, and merely described rather than perceived. When someone reads over a description of an agent whose perceptual sensitivity to a range of environmental affordances is apt to be integrated with her practical reasoning, the connection of such an account to her own perceptual experience can seem opaque – according to an action-space-inspired version of PCS, this is because the linguistically-mediated and conceptual manner in which the reader conceives of the affordances in question is very different from the engaged and preconceptual grasp of currently salient affordances that is crucial for conscious experience according to the action-space approach¹⁶⁰.

¹⁶⁰ On this account, the support the action-space affords to PCS is thus slightly different from that which it affords access-based and representationalist strategies. The latter strategies

These remarks on how the action-space approach as we have developed it supports the three proposals allow us to see what they have in common. Each attempts to deal with the explanatory gap by claiming that we are prone to making a certain sort of mistake when, by reflecting on our experience, we fix the concept of phenomenal consciousness we're trying to explain. Clark suggests that we mistake our limited access to the features in virtue of which our sensory episodes differ for confrontation with some non-structurally or functionally analyzable explanandum. Jackson and the representationalists say that we mistake properties of the way our experience represents things as being for instantiated properties of our experience. And phenomenal concept strategists say that we mistake the conceptual isolation of the phenomenal concepts we use to think about experience for a kind of metaphysical isolation of the phenomenal from the structural and functional. Viewing the proposals in this way suggests how I think the action-space theorist should respond to the explanatory gap.

In brief, the action-space theorist holds that conscious experience consists in being put in touch with the affordances of your environment in a way that enables intelligent action, but that certain facts about our mode of access to those affordances make it natural to think we are confronted in experience with phenomenal properties that go beyond mere specifications of possibilities for action. In slightly more detail, the action-space theorist agrees with the phenomenal concept strategist that the appearance of an unbridgeable explanatory gap between the phenomenal and the physical arises since we are prone to thinking about conscious experience in two distinct ways. One of these ways (the one that gives rise to the apparent explanatory gap) relies on a distortive conception of experience – as Jackson and the representationalists suggest, we are prone to reify facts about how our experience represents the world and our relation to it into putative facts about phenomenal particulars with which we are confronted in experience. Specifically, we are prone to move from the fact that an experience manifestly empowers us to act in certain ways that can be reflected in our practical reasoning to the putative fact that there is some further instantiated phenomenal property of experience that stands in need of explanation. But we can give an account of why this distortive conception of experience is natural – as Clark suggests, the access we have to the skills enabled for us by our perceptual sensitivity is subject to certain limits. We suggested

are supported by the arguments in favour of the action-space account's conception of visual experience in chapters 2, 3 and 4. I am suggesting that PCS can be fleshed out via an appeal to the model of practical reasoning and grasp of enabled actions suggested in chapter 5.

above that these limits are imposed by the discrepancies between the preconceptual grasp of enabled actions necessary for the form of practical reasoning outlined in chapter 5, and the linguistic, conceptual grasp we have of the affordances that correctly characterise the content of perceptual experience when they are presented to us as a verbal description of the conditions that suffice for conscious experience. Thus, when we reflect on our experience, and attempt to describe the ways in which two visual experiences differ, the full range of differences between the suites of enabled actions that characterise each experience is not accessible to verbal report or conceptualisation. It thus does not seem to us, on the basis of introspective reflection, that experiences differ only in virtue of the way they poise us over a space of enabled actions. Once we have explained this in terms of the distinctive features of the access we have to such a space of enabled actions, we see the possibility that the obtaining of such knowing poise and patterns of access can ground a complete account of phenomenal consciousness. It might still be tempting for some to suppose that the relevant patterns of access involve acquaintance with essentially experiential phenomenal particulars, but PCS and the representationalists have shown us how such an inference may rest on a confusion, and Clark's demonstration of how facts about access can entail facts about how things seem from a particular perspective make such an inference less tempting still.

As mentioned above, I think that viewing the three proposals through the lens of the action-space theory not only discloses the relations between them, but suggests how each might respond to the second class of questions I raised for them in the previous section. I asked whether we should read Clark's proposal as an attempt to build a conceptual bridge between patterns of access to differences between sensory episodes and phenomenal consciousness, or to assuage the need for such a conceptual bridge. Our discussion here suggests that we should understand the proposal as offering a conceptual link between phenomenal consciousness *properly construed* (as appropriate grasp of the practical consequences of perceptual sensitivity) and patterns of access, as well as an account of why such a construal of perceptual experience is not introspectively obvious. We considered a development of Chalmers' objection, where it was claimed that Clark's proposal can explain propensities to judge or report that one is having conscious experiences, but leaves out the phenomenal truthmakers of such judgements and reports. In response, we should hold that the intuition that some phenomenal property that is in principle independent of the ways the subject is disposed to act and practically reason stands behind the subject's introspective judgement of difference between two experiences stems from a natural but distortive conception of experience which we have been

given reason to think is a mere artefact of the discrepancy between the limited but intimate access to environmental affordances we have in experience, and the linguistic and conceptually-mediated access we have to them as they figure in our description of the sufficient conditions for conscious experience.

Similarly, we should see representationalists as suggesting a conception of experience according to which there can be a conceptual mesh between conscious experience, properly characterised, and a physicalistically-acceptable explanans (understood, on our account as a certain way of representing the current possibilities for action afforded by one's environment). I asked how representationalists might explain residual intuitions that the representational contents to which they appeal might be tokened in the absence of phenomenal consciousness. The representationalist can invoke our discussion of Clark and PCS in response. We have suggested that there are different modes of access to contents involved when actively perceptually engaged with one's environment on the one hand, and indulging in disengaged conceptual reflection on those contents on the other, and that these differences can explain the apparent gulf between the ways experiences strike us introspectively, and the structural and functional characterisations we can give of those experiences. However, Clark's proposal also makes it plausible to suppose that when patterns of access of the kind to which we appealed in chapter 5 and above obtain, an agent will sincerely make judgements on the basis of introspection such as "there is 'something it's like' to undergo this experience as opposed to that one", clarifying how a representational state can bring about things feeling a certain way for an agent. Note also the added resources to explain the temptation to the represented/instantiated confusion afforded by viewing the representationalist strategy in light of the proposals of Clark and PCS.

Lastly, I questioned whether PCS could do more than demonstrate the compatibility of our inability to construct a naturalistically-acceptable explanation of conscious experience with the truth of physicalism. I also noted (in fn. 13, following Carruthers and Veillet (2007)) that the truth of PCS is not incompatible with conscious experience being subserved by non-physical properties. As Carruthers and Veillet note, the way we have sketched PCS relies on a distinction between 1st and 3rd-person characterisations of, or ways of picking out, phenomenal concepts. 1st-person characterisations are those that presuppose a particular perspective, such as "the concept I deploy when thinking about *this* experience", or "the recognitional concept which applies to states like *this* and *that*." A 3rd-person characterisation such as "phenomenal

concepts are conceptually isolated recognitional concepts that are applicable only in the presence of the states which they pick out” presupposes no such perspective. The 1st-person characterisations are intended to illustrate the way phenomenal concepts figure in our thoughts about our experiences, whilst the 3rd-person characterisations are intended to open up the possibility that our 1st-person characterisations can be given a reductive explanation. However, we might worry that we have affirmed the existence of an explanatory gap, but relocated it between our 1st-person and 3rd-person accounts of phenomenal concepts. We noted above that an anti-physicalist might hold that a distinctive kind of phenomenal concept does figure in our 1st-person thoughts about experience, and serves to pick out intrinsic, non-physical phenomenal properties. If PCS cannot provide a story about the relations between the 1st-person construals of phenomenal concepts and their physicalistically-acceptable 3rd-person construals, the anti-physicalist can claim that the absence of a conceptual mesh between the 1st and 3rd-person characterisations points to the existence of an explanatory gap between our 1st-person experience and the 3rd-person structural and functional vocabulary with which we wish to explain it.

But our consideration of Clark and the representationalists suggests how PCS should respond here. If we hold that, as the action-space suggests, the apparent discrepancy between 1st and 3rd-person methods of conceiving of our experience can be explained by appeal to the discrepancies between the preconceptual grasp of environmental affordances necessary for conscious experience and the conceptual grasp of those affordances as they figure in a 3rd-person characterisation of the content and character of an experience then Clark’s proposal has shown us how patterns of access to afforded access, characterised in 3rd-person terms, can entail there being something-it’s-like for the agent for whom the actions are afforded. Whilst it is still logically possible that the relevant patterns of access involve acquaintance with non-physically-analysable explananda, the suggestions of Jackson and the representationalists have made it plausible that drawing this conclusion on the basis of our introspective evidence rests on a confusion, and the suggestion that Clark’s proposal shows how facts about access could entail facts about what-it’s-like from a particular perspective removes what I take to be the most plausible motivation for embracing such a possibility.

It seems, then, that viewing our proposals through the lens of the action-space theory has suggested ways in which each might be further developed and empirically-supported, as well as how the proposals share a common structure that allow them to work together in order to resist the criticisms that arise for each. I want to close by clarifying how our discussion of the

three proposals and the action-space theory in this section suggests a unified solution to the explanatory gap that meets the conditions on an adequate deflationist theory we set out in 1.4.

6.4 The Duties of the Deflationist Revisited

In 1.3 I suggested that we view the task of constructing a deflationist explanation of conscious experience as analogous to explaining a magic trick. On a natural interpretation of the trick, we are confronted with events that resist explanation in terms of our standard physical ontology, and we may be driven to various anti-reductionist theories as a result. However, empirical and conceptual investigation into the trick lets us see that a description of the trick that makes ineliminable references to vanishings, teleportations and suchlike is not the only way to accommodate our evidence – we can instead explain it via an appeal to sleight-of-hand, a phenomenon that can be explained in physicalist terms, whilst explaining away the temptation to construe the trick as involving genuine magic. Last section's discussion of the action-space account and the three proposals suggests a similar position with regard to conscious experience. Having a conscious experience with a certain content and character turns out to be a matter of grasping the ways one's perceptual sensitivity to one's environment currently poises one to act. However, facts about the nature and limitations of our access to the suites of affordances that contribute to the content of our conscious experience mean that this fact is not introspectively obvious to us – rather, we are prone to suppose that we are acquainted in experience with phenomenal particulars that are essentially neutral with respect to their relations to possible actions on the part of the agent¹⁶¹. Both this type of poise over a space of actions, and the limits of our access to the nature and extent of those actions are amenable to reductive explanation. It seems to me that the account of phenomenal consciousness we have arrived at thus meets the desiderata we set out in 1.4 of accommodating the introspective evidence that leads to the positing of an explanatory gap, but redescribing that evidence in such a way as to make it compatible with physicalism, and showing why the physicalistically-acceptable construal of that evidence is not introspectively obvious.

¹⁶¹ An interesting and open question arises about the extent to which the temptation to think about experience in this way is universal, culturally informed, or the product of studying a certain sort of philosophy. The difficulties I sometimes have explaining to people without a background in philosophy why there is thought to be a particular problem about the reductive explanation of consciousness makes me curious about the latter possibility, but I shan't attempt to speculate further here.

In general, I have stressed the obligation of a theory of consciousness to say something about how it meets (or at least relates to) the requirement that reductive explanations demonstrate a conceptual mesh between explanandum and explanans. According to our action-space account, no such conceptual mesh is in the offing on one construal of our introspective evidence. This prompts the conclusion that there is an unbridgeable explanatory gap between the phenomenal and the physical. But our discussion of the three positions above has shown us how to resist the conclusion – the construal of our introspective evidence required to motivate the explanatory gap is natural, but distortive. We argued that we can also view our introspective evidence as consisting in our automatic and non-inferential practical grasp of sets of possibilities for various sorts of action suggested by our current perceptual sensitivity to the environment – experience construed in this way is compatible with a physicalist explanation that includes an *a priori* component¹⁶².

Additionally, the arguments in favour of the action-space conception of experience in chapters 2, 3 and 4, and the model of practical reasoning and access to perceived affordances in chapter 5, suggest that our redescription is empirically plausible, and thus that the conditions the action-space account specifies for phenomenal consciousness are indeed met by conscious perceivers.

We also noted in 1.4 that a candidate deflationist account of phenomenal consciousness should say something about the properties of conscious experience that are sometimes held to make it a problematic candidate for reductive explanation. We listed ineffability, privacy, direct-accessibility and intrinsicality as examples, and concluded that the deflationist should either explain how such properties can arise as a result of physicalistically-acceptable conditions for conscious experience obtaining, or explain how our introspective evidence makes it natural to posit such properties despite the fact that no such properties of experience obtain. Our discussion above suggests how we might discharge this obligation for each property.

We can begin to explain the ineffability of conscious experience by noting the complexity of the enabled abilities that characterise visual experience. For example, our discussion of colour experience in chapter 4 suggested that experiencing a surface under natural light as a uniform colour is a matter of being able to sift, sort and track it according to its location in a multi-dimensional space of possible ways of modifying light, that this ability in turn was enabled by grasping the significance and likely patterns of variations in the array of

¹⁶² Recall that I have assumed that representational and intentional properties can be reductively explained in physical terms.

colour-looks it displayed, and that perceptual sensitivity to those colour-looks in turn was a matter of being able to place different portions of the surface in a multi-dimensional space whose geometry is dictated by relations of relative sameness and difference. The nature and extent of the abilities which form the basis of a standard visual experience of colour, then, are difficult both to catalogue and to convey. Moreover, our access to those enabled abilities is limited by the possible constraints we set out in chapter 5 – the practical grasp of enabled abilities necessary for conscious experience according to our account does not entail the abilities to conceptualise and report on those abilities. We brought this point out in our above discussion of Clark’s creature with limited non-inferential access to features in virtue of which its sensory episodes differ. We would expect such a creature to find its sensory experiences ineffable, since the nature of its access to the properties of those experiences imposes limits on the detail in which it can conceptualise and report upon them.

An account of privacy and direct availability is also suggested by such considerations about access. The experiences of a creature with the patterns of access to which we have been appealing will always be private, firstly since it is unable to fully convey the character of its experiences in its reports to others, and secondly since, for any conscious sensory episode, it alone enjoys the access to the range of affordances disclosed for it by its perceptual sensitivity to the environment that suffices for conscious experience according to the action-space account. Phenomenal properties of experiences are likewise directly-available to such a creature since a detected affordance only figures in an agent’s experience when it is apt to inform her ongoing practical reasoning, thus ruling out the possibility that there could be phenomenal properties of the agent’s experience that were inaccessible to her¹⁶³.

Lastly, the considerations of the previous two sections, in particular our discussion of representationalism, suggest how the action-space account should treat the putative experiential property of intrinsicity. We noted in chapter 1 that we cannot think of phenomenal properties as ‘intrinsic’ if we hope to construct a reductive explanation of consciousness. This was because doing so seemed to entail the essential independence of those properties from any structurally or functionally specifiable corollaries. However, as suggested by our quote from Rorty in 6.2.2, such a conception of experience shifts from viewing experiences as essentially properties of persons to viewing them as phenomenal particulars whose essence is exhausted by the way they appear to experiencing subjects. As we saw there (and in 1.4), such a conception is incompatible with physicalism about experience. However, we also saw

¹⁶³ So long as we are construing aptitude for integration with practical reasoning as a form of availability or accessibility. This is in line with Block’s (1995) standard definition of ‘access’.

as a result of our discussion of representationalism and its relation to the other positions of 6.2 that we have reason to think that such a conception of experience is misguided. Viewing experiences as involving confrontations with instantiated intrinsic phenomenal properties involves the equivocation Jackson highlighted between instantiated and intentional properties of experience. In the last section we saw how the facts about access to which the action-space account appeals suggests that such an equivocation is tempting but non-compulsory, and we have good reason to think it distortive. Intrinsicity, then, is at best an apparent property of conscious experiences.

Our action-space brand of deflationism, then, meets the desiderata we set out on an adequate explanation of conscious experience, addressing the conceptual mesh requirement, giving a plausible and physicalistically-acceptable construal of our introspective evidence whilst explaining the temptation to an alternative construal, providing empirical evidence in its favour, and affording plausible diagnoses of each of the putative problematic properties of experience we have considered.

Conclusion

I have argued in this chapter that the action-space account of conscious vision allows us to see how to close the explanatory gap. We need not view the action-space account as merely a proposal about the contents of visual experience, a thesis about the functional role of conscious experience, or a way of addressing the comparative explanatory gap, though it provides all these things. We saw above, though, that it can also make intelligible how facts about an agent's access to the skills enabled for them by their perceptual sensitivity to their environment can entail facts about how things seem from that agent's perspective, can result in ineffability, privacy and direct-availability being properties of that agent's experience, and make it natural (though mistaken) for that agent to suppose that she is confronted in her experience with intrinsic phenomenal properties which resist structural and functional explanation. Drawing on Clark's access-based strategy, representationalism and PCS, we traced the source of this confusion to facts about the distinctive mode of access an agent has to the skills enabled for her when those skills form the basis of a conscious experience. The preconceptual, practical and immediate accessibility of those skills to the agent for whom they are manifestly enabled is unlike the conceptual, theoretical and inferential access to

those skills as they figure in a verbal characterisation of the content of that agent's experience, creating the impression of an unbridgeable gap between 1st and 3rd-person characterisations of that experience. However, Clark's proposal about access suggested that facts about how things are from a 1st-person perspective can follow from 3rd-personal facts about patterns of access, and our consideration of representationalism and PCS suggested that we can resist moving from the residual intuition that such a conception leaves features of experience unexplained to anti-physicalist conclusions. Instead, I suggested that we explain this intuition away as a confusion between intentional and instantiated properties of our experience that arises due to facts about the nature of our access to those intentional properties in conscious perception. Cashing out the three proposals of 6.2 in action-space terms enabled us to see their common structure, to suggest how they can work together to accommodate the criticisms faced by each and, thus unified, benefit from the empirical support for the action-space theory provided in chapters 2, 3 and 4. We emerge, I think, with good reason to believe that conscious vision, and not merely its perceptual content or functional role, can be explained in terms of an agent's grasp of the practical consequences of their perceptual sensitivity – in terms of knowing poise over a space of actions.

Conclusion

Conscious visual experience, I have argued, consists in an agent's exercising a certain kind of grasp of the practical consequences of their perceptual sensitivity to their environment. Episodes of conscious visual experience occur when an agent's perceptual sensitivity to the mesh between their current environment and their current goals and projects is apt to inform their practical reasoning. Evidence in favour of such a view has come from considering what the DVS data tells us about the functional role of conscious vision (ch.2), considering how such an account can handle problem cases which confound other enactive theories and present problems for any candidate account of perceptual consciousness (ch.3), and the merits of the account of colour perception which it suggests (ch.4). I suggested that the requirement that perceptual sensitivity to affordances be "apt for integration with practical reasoning" if it is to issue in conscious experience should be understood as requiring that the perceiver grasps both that the perceived afforded action is a means to one among several possible ends, and the relevance of such means/ends relationships for her current goals. We also considered the possibility that such a view suggests how questions about the perspectival nature and qualitative character of conscious experience are tackled together for the action-space approach – the conditions that bring about experiences with qualitative character might also suffice for the existence of a unified, normatively constrained perspective (ch.5). Lastly, I argued that the action-space account lets us see how the belief that there is an unbridgeable explanatory gap between the phenomenal and the physical rests on a natural but distortive perception of experience, and shows us how to remedy the weaknesses and combine the strengths of existing attempts to address the explanatory gap (ch.6).

I'd like to finish by noting some areas where the action-space account I have sketched might be improved, and some possible avenues for further research. Firstly, much more needs to be said about the specific sorts of actions and propensities for worldly engagement to which the action-space account appeals. I focussed in some detail on colour experience in chapters 3 and 4, suggesting that in that perceptual domain abilities to sift, sort, compare, contrast and track can ground an adequate account of the content and character of experience. But my focus on colour experience was motivated in large part by the fact that it appears to be an atypical case for the action-space account – a class of perceptual experience where the qualitative dimension of the experience seems intuitively independent of the sorts of actions that might be associated with it. We should thus investigate the possibility that other domains of visual

experience might possess their content and character in virtue of very different spaces of actions. Rick Grush's (1998, 2007) work on the behavioural basis of spatial perception, or Matthen's (2005) discussion of the visual experience of presence both suggest how propensities to act of a different kind from those which we have mostly been concerned might contribute to an account of aspects of visual content. The action-space account would benefit from investigating whether any unified account of the actions which form the bases of episodes of conscious vision can be given in light of work such as this.

Secondly, more remains to be said about the relationship between the action-space and sensorimotor views. For example, Hurley (1998, ch.6) warns against the "Myth of the Giving" (the corollary to the discredited "Myth of the Given" concerning perceptual content) – according to which the content of intentions and actions is taken as basic and unproblematic. Hurley's claim is not that this thesis is often explicitly endorsed, but that work in the philosophy of mind often helps itself to the assumption that some intention or action has a specific content in a way that Sellars (1997), McDowell (1996) and others have rendered unacceptable for the content of perceptual episodes, and that there is no compelling theoretical justification for this asymmetry. Hurley's desire to avoid both the Myth of the Given and the Giving is one of the considerations that leads her to a view of perception and action as dynamically interdependent, whereby the content of a perceptual state is fixed by its implications for possible actions, and the content of an intentional action is fixed by its implications for subsequent perceptions. Hurley's "sensorimotor" view thus combines the action-space account's emphasis on the practical consequences of perception with Noë's emphasis on the practical consequences of action. The action-space theorist needs to consider whether they have succumbed to the Myth of the Giving – certainly, in our account here, we have argued that the content and character of perceptual experience should be accounted for in terms of the perceiver's current grasp of available opportunities for intentional action, without giving an account of what distinguishes the content of one enabled action from another. Unanswered questions surface here, then, as to whether the action-space theorist is obliged to give an account of intentional content, whether they can do so without their view collapsing into Hurley's dynamic interdependence theory, and whether such a collapse would undermine aspects of the account we have developed here.

There are many other points in the action-space account as I have sketched it here that would benefit from further discussion and elaboration, and I shan't try and enumerate them here.

The action-space account of consciousness, we should admit, is perhaps not yet best viewed as a full-fledged theory, but as a research programme centred around the idea that conscious experience can be fruitfully understood in terms of a grasp of the practical consequences of perceptual sensitivity. This thesis has suggested how I think this guiding idea can be supported and fleshed out, and I want to close by suggesting how it might apply in domains other than those already considered. One interesting avenue for future research concerns the extent to which the action-space account we have sketched might be generalised to other aspects of conscious experience. Our arguments have been largely confined to showing that visual sensitivity to the affordances of our environment issues in conscious experience when apt for integration with capacities for practical reasoning, but there are several ways in which we might attempt to extend our framework to other perceptual modalities. Firstly, we have used the dual visual systems data to support our claim that visual information that carries implications about the current opportunities for the satisfaction of an agent's intentional plans and projects, rather than information apt to specify the details of the movements necessary to achieve those plans and projects, makes up the content of conscious visual experience. There is evidence that a similar duality of processing obtains for hearing as well as vision. A ventral stream appears to carry information about the enduring semantic properties of the source of a sound, whilst a dorsal stream carries detailed information about the source of that sound relative to the perceiver (Arnott et al (2004), Rauschecker and Tian (2000)). Evidence for such a model is still emerging, but it will be interesting to see how closely the parallels between functional profiles of auditory and visual streams run. If those parallels extend to the features of ventral processing we identified as relevant to the action-space account in chapters 2 and 3, then a version of the action-space account generalised to auditory perception would predict illusions and pathologies of aural experience analogous to those discussed in chapter 2. Cases of "deaf hearing" (Mozaz-Garde and Cowey (2000)) that seem to be auditory analogues of blindsight have indeed been reported – it might be that the neuropsychology behind such cases makes it plausible to diagnose them in an analogous way to the action-space account's treatment of DF and blindsight in chapters 2, 3 and 5. Similarly, the plausibility of an action-space account of tactile experience might be investigated by considering the possibility of an action-space account of "numbsense", an apparent somatosensory analogue of blindsight (Rossetti et al (1995)).

More generally, the action-space account of colour perception sketched in chapters 3 and 4 suggests how the action-space approach might be extended to sensory perception more

generally. Recall that a keystone of our strategy there was to place the colours to which an agent was perceptually sensitive in a space whose geometry is dictated by relations of relative similarity and difference. We then claimed that an agent's perception of a particular colour consisted in their manifesting a practical grasp of the way a visually-presented object or region stands in a complex set of relations of similarity and difference to the other colours to which they are perceptually sensitive. In summary, we constructed for the domain of colour qualities to which perceivers are visually sensitive a relationally-defined quality space, and then argued that such a quality space obtaining for a perceiver consists in their possession of a certain network of practical abilities. Our construction of this account made much use of conceptual apparatus for understanding sensory perception in general that has been developed in detail by Austen Clark (1993) and Matthen (2005). Clark sees the ordering of sensory qualities into quality spaces whose geometries are defined by relations of perceptible similarity and difference as a general approach to sensory perception that extends beyond the domain of colour, and Matthen argues for a view of the senses as automatic sorting systems that class perceived stimuli according to relations of graded similarity and issue in information about how stimuli should be treated in light of the overall goals of the perceiver. However, both Matthen and Clark focus chiefly on the modality of vision, and the specific case of colour, in order to illustrate their conception of sensory perception. A further avenue for research, then, is investigating the extent to which such approaches are supported by current empirical work on perception in other modalities, and whether the accounts they might provide are also apt to provide support for the action-space account.

Lastly, we might investigate whether the action-space account affords the materials for an account of our experiences of emotions. Whilst emotions seem like a prime example of private subjective states that can often be suppressed, rather than manifested, in overt behaviour, there is nonetheless empirical and philosophical precedent for viewing emotions as inextricably tied to the ways we are disposed to act in the world. Emotional experience, we might argue, consists in a sense of the way in which currently-obtaining possibilities for action are restricted, or fail to be so restricted. The possibilities for action suggested to a perceiver in a state of joyous enthusiasm might be very different from those suggested by the same scene to a perceiver in the depths of despair. In this vein, Radcliffe (2008), following Heidegger, argues that the experience of mood should be understood in terms of the notion of "possibility-spaces", and Frijda (2006) emphasises the essential motivational components of affective states in his account of the emotions.

However, the challenge that is both most formidable and most exciting for the action-space approach goes beyond the above questions of how the approach might apply to other perceptual modalities. It concerns the vision of mind that lies at the heart of the action-space approach – as always and everywhere environmentally embedded, engaged and active. What does such a view of the mind tell us about the relation between mind and world, and how might such a view reconfigure traditional problems in the philosophy of mind? In chapter 5 I examined, following Hurley, how an enactive focus might inform our view of practical reasoning and deliberation, and suggested that the tasks of explaining self-consciousness and of explaining phenomenal consciousness might coincide for an enactive theory. In chapter 6, I argued that the action-space account lets us see how a traditional philosophical conception of experience may be misguided, and how the view of experience suggested by the action-space account affords a solution to the problem of the explanatory gap. These discussions merely hint at the beginnings of the ways in which viewing the mind as essentially active, engaged and environmentally intertwined can shed new light and afford new perspectives on existing philosophical discussions. The most exhilarating thing about the enactive view of the conscious mind is the prospect that philosophy might look very different once its ramifications are truly understood.

Bibliography

- Aglioti, S., DeSouza, J.F.X. and Goodale, M.A. (1995) 'Size Contrast Illusions Deceive the Eye but not the Hand', *Current Biology*, 5(6), 679-685
- Arnott, S, Binns, M.A, Grady, C.L and Alain, C. (2004) 'Assessing the Auditory Dual-pathway Model in Humans', *Neuroimage*, 22(1), 401-408
- Aydede, M. and Guzeldere, G. (2005), 'Cognitive Architecture, Concepts and Introspection', *Nous*, 39(2), 197-255
- Bach-y-Rita, P. and Kercel, S.W. (2003) 'Sensory Substitution and the Human-machine Interface', *Trends in Cognitive Sciences*, 7(12), 541-546
- Beck, J. (ms.) 'The Generality Constraint, Analog Magnitudes and the Structure of Cognition'
- Berlin, B. and Kay, P. (1969) *Basic Colour Terms: Their Universality and Evolution*, Berkeley, University of California Press
- Bermudez, J.L. (1998) *The Paradox of Self-Consciousness*, Cambridge, MIT Press
- _____ (2001) 'Nonconceptual Self-Consciousness and Cognitive Science', *Synthese*, 129, 129-142
- Bermudez, J.L. and MacPherson, F. (1998) 'Nonconceptual content and the nature of perceptual experience' *The Electronic Journal of Analytic Philosophy*, 6.
- Block, N. (1995) 'On a Confusion About a Function of Consciousness', *Behavioural and Brain Sciences*, 18(2), 227-247
- _____ (2005) 'Review of Alva Noë, *Action in Perception*', *Journal of Philosophy*, CII(5), 259-272
- _____ (2007) 'Consciousness, Accessibility, and the Mesh Between Psychology and Neuroscience', *Behavioural and Brain Sciences*, 30, 481-548
- Block, N. and Stalnaker, R. (1999) 'Conceptual Analysis, Dualism and the Explanatory Gap', *Philosophical Review*, 108, 1-46
- Brenner, E. and Smeets, J.B.J. 'Size Illusion Influences how we Lift but not how we Grasp an Object', *Experimental Brain Research*, 111, 473-476
- Broackes, J. (1992) 'The Autonomy of Colour', in Charles and Lennon (ed) (1992) *Reductive Explanation and Realism*, Oxford, Oxford University Press
- _____ (2003) 'Do Opponent-process Theories Help Physicalism about Colour?', *Behavioural and Brain Sciences*, 26(6), 786-788
- _____ (2007) 'Black and White and the Inverted Spectrum', *The Philosophical Quarterly*, 57, 161-175

- Carruthers, P. (2000) *Phenomenal Consciousness*, Cambridge University Press
- _____ (2004) 'Reductive Explanation and the "Explanatory Gap"', *Canadian Journal of Philosophy*, 34, 153-173
- Carruthers, P. and Veillet, B. (2007) 'The Phenomenal Concept Strategy', *Journal of Consciousness Studies*, 14(9), 212-236
- Cartwright, R. (2004) 'Sleepwalking Violence: A Sleep Disorder, A Legal Dilemma and a Psychological Challenge', *American Journal of Psychiatry*, 161, 1149-1158.
- Chabris and Simons (1999) 'Gorillas in our Midst: Sustained Inattentional Blindness for Dynamic Events', *Perception*, 28, 1059-1074
- Chalmers, D. (1996) *The Conscious Mind*, Oxford, Oxford University Press
- _____ (1997) 'Facing up to the Problem of Consciousness', in Shear, J. (ed.) *Explaining Consciousness: The Hard Problem*, Cambridge, MIT Press
- _____ (2003) 'Consciousness and Its Place in Nature', in S. Stich and F. Warfield (eds.), *Blackwell Guide to Philosophy of Mind*, Blackwell
- _____ (2004) 'How Can We Construct a Science of Consciousness?' in (M. Gazzaniga, ed) *The Cognitive Neurosciences III*. MIT Press
- Chalmers, D. and Jackson, F. (2001) 'Conceptual Analysis and Reductive Explanation', *Philosophical Review*, 110, 315-360
- Chartrand, T. and Bargh, J. (1999) 'The Chameleon Effect: The Perception-Behaviour link and Social Interaction', *Journal of Personality and Social Psychology*, 76(6), 893-910
- Churchland, P. (1997) *The Engine of Reason, The Seat of the Soul*, Cambridge, MIT Press
- _____ (2005) 'Chimerical Colours; Phenomenological Predictions from Cognitive Neuroscience', *Philosophical Psychology*, 18(5), 527-560
- Clark, Andy (2000a) 'A Case Where Access Implies Qualia?', *Analysis*, 60(265), 30-38
- _____ (2000b) *Mindware: An Introduction to the Philosophy of Cognitive Science*, USA, Oxford University Press
- _____ (2001) 'Visual Experience and Motor Action: Are the Bonds Too Tight?', *Philosophical Review*, 110, 495-519
- _____ (2006a) 'Vision as Dance: Three Challenges for Sensorimotor Theory', *PSYCHE*
- _____ (2006b) 'Cognitive Complexity and the Sensorimotor Frontier', *Proceedings of the Aristotelian Society*, Supplemental Volume, 80(1), 43-65
- _____ (2007) 'What Reaching Teaches: Consciousness, Control and the Inner Zombie', *British Journal for the Philosophy of Science*, 58, 563-594

- Clark, Andy and Grush, R. (1999) 'Towards a Cognitive Robotics', *Adaptive Behaviour*, 7(1), 5-16
- Clark, Andy and Toribio, J. (2001) 'Sensorimotor Chauvinism?', *Behavioural and Brain Sciences*, 24, 979-980
- Clark, Austen (1993) *Sensory Qualities*, Oxford, Clarendon Press
- _____ (2004) *A Theory of Sentience*, Oxford, Oxford University Press
- _____ (2007) 'Color, Qualia and Attention', in Matthen, M. and Cohen, J. (eds) (2007) *Color Ontology and Color Science*, Cambridge, MIT Press
- Cohen, J. (2000) 'Color: A Functionalist Proposal' *Philosophical Studies*, 113, 1-42
- Cowey, A. (2004) 'Fact, Artefact and Myth about Blindsight' *The Quarterly Journal of Experimental Philosophy, Section A*, 57(4), 577-609
- Crane, T. (2000) 'The Origins of Qualia', in Crane, T. and Patterson, S. (eds) *The History of the Mind-Body Problem*, London, Routledge
- Crisp, A.H., Matthews, B.M., Oakey, M., Crutchfield, M. (1990) 'Sleepwalking, night terrors and consciousness' *British Medical Journal*, 300, 360-62.
- Daprati, E. and Gentilucci, M. (1997) 'Grasping an Illusion', *Neuropsychologia*, 35(12), 1577-1582
- Davies, M. (1992) 'Aunty's own argument for the language of thought' In J. Ezquerro and J.M. Larrazabal (eds.), *Cognition, Semantics and Philosophy: Proceedings of the First International Colloquium on Cognitive Science*. Dordrecht, Kluwer Academic Publishers, 235-71.
- Dennett, D. (1998) 'Quining Qualia', in Marcel and Bisiach (eds.) *Consciousness in Contemporary Science*, Oxford, Oxford University Press
- _____ (1991) *Consciousness Explained*, Boston, Little Brown
- _____ (1993) 'Living on the Edge', *Inquiry*, 36
- Dretske, F. (2006) 'Perception Without Awareness', in Gendler, T.S. and Hawthorne, J., *Perceptual Experience*, Oxford University Press
- Edquist, J., Rich, A.N., Brinkman, C. and Mattingley, J.B. (2006) 'Do Synaesthetic Colours Act as Unique Features in Visual Search?', *Cortex*, 42, 222-231
- Ellis, R.R., Flanagan, J.R. and Lederman, S.J. (1999) 'The Influence of Visual Illusions on Grasp Position', *Experimental Brain Research*, 125, 109-14
- Evans, G. (1982) *The Varieties of Reference*, edited by John McDowell, Oxford, Oxford University Press
- Fodor, J. (1974) 'Special Sciences or: The Disunity of Science as a Working Hypothesis', *Synthese*, 28, 97-115.

- _____ (1987) *Psychosemantics*, Cambridge, MIT Press
- Fodor, J. and Lepore, E. (1992) *Holism: A Shopper's Guide*, Blackwell Publishing
- Foster, J. (1989) 'A defense of dualism', in Smythies, J. and Beloff, J. (eds), *The Case for Dualism*, University of Virginia Press.
- Frijda, N. (2006) *The Laws of Emotion*, Lawrence Erlbaum Associates
- Frith, C. (1992) *The Cognitive Neuropsychology of Schizophrenia*, Psychology Press
- Gallistel, C.R. (1990) *The Organization of Learning*, Cambridge MA, Bradford books, MIT Press
- Ganel, T. and Goodale, M. (2003), 'Visual Control of Action but not Perception Requires Analytical Processing of Object Shape', *Nature*, 426, 664-667
- Grush, R. (1998) 'Skill and Spatial Content', *The Electronic Journal of Analytic Philosophy*, 6
- _____ (2007) 'Skill Theory v.2.0: Dispositions, Emulation and Spatial Perception', *Synthese*, 159(3), 389-416
- Goodale, M. and Milner, A.D. (2004) *Sight Unseen: An Exploration of Conscious and Unconscious Vision*, Oxford University Press
- Grossenbacher, P.G. and Lovelace, C.T. (2001) 'Mechanisms of Synaesthesia: Cognitive and Physiological Constraints', *Trends in Cognitive Sciences*, 5(1), 36-42
- Haffenden, A.M. and Goodale, M. (2000) 'Independent Effects of Pictorial Displays on Perception and Action', *Vision Research*, 40, 1597-1607
- Haffenden, A.M., Schiff, K.C. and Goodale, M.A. (2001) 'The Dissociation Between Perception and Action in the Ebbinghaus Illusion: Non-illusory Effects of Pictorial Cues on Grasp', *Current Biology*, 11, 177-181
- Hardin, C.L. (1988) *Color for Philosophers: Unweaving the Rainbow*, Indianapolis, Hackett.
- Harman, G. (1990) 'The Intrinsic Quality of Experience', in Tomberlin, J. (ed.) *Philosophical Perspectives*
- Haugeland, J. (1990) 'The Intentionality All-Stars', in Tomberlin, J. (ed.) *Philosophical Perspectives*
- Hill, C. (1997) 'Imaginability, Conceivability, Possibility and the Mind-body Problem', *Philosophical Studies*, 87(1), 61-85
- Hurley, S. (1998) *Consciousness in Action*, Harvard, MIT Press
- _____ (2001) 'Reason and Motivation: The Wrong Distinction?' *Analysis*, 61(270), 151-155
- _____ (2004) 'Active Perception and Perceiving Action: The Shared Circuits Model', in Gendler, T.S. and Hawthorne, J., *Perceptual Experience*, Oxford University Press

- _____ (2006) 'Making Sense of Animals', in Hurley, S. and Nudds, M. (eds) *Rational Animals?*, NY, Oxford University Press
- _____ (2008) 'The Shared Circuits Model (SCM): How Control, Mirroring and Simulation can Enable Imitation, Deliberation and Mindreading', *Behavioural and Brain Sciences*, 31, 1-22
- _____ (forthcoming) 'Varieties of Externalism', in Menary, R. (ed) *The Extended Mind*, Aldershot, Ashgate
- Hurley, S. and Noë, A. (2003) 'Neural Plasticity and Consciousness', *Biology and Philosophy*, 18(1), 131-168
- _____ (2007) 'Can Hunter-gatherers Hear Colour?' in Brennan, G. (ed) *Common Minds: Essays in Honour of Philip Pettit*, USA, Oxford University Press
- Hurvich, L.M. and Jameson, D (1957), 'An Opponent-process Theory of Colour Vision', *Psychological Review*, 64(6,pt.1), 384-404
- Ingle, D.J. (1973) 'Two Visual Systems in the Frog', *Science*, 181, 1053-1055
- Jackson, F. (1982) 'Epiphenomenal Qualia', *The Philosophical Quarterly*, 32, 127-136
- _____ (2003) 'Mind and Illusion', In O'Hear, A. (ed.), *Minds and Persons*, Cambridge University Press
- Jackson, S.R. and Shaw, A. (2000) 'The Ponzo Illusion Affects Grip Force but not Grip-aperture Scaling During Prehension Movements', *Journal of Experimental Psychology*, 26, 418-423
- Jacob, P. (2006) 'Why Visual Experience is Likely to Resist Being Enacted', *PSYCHE*
- Jacob, P. and Jeannerod, M. (2003) *Ways of Seeing: The Scope and Limits of Visual Cognition*, NY, Oxford University Press
- James, T.W., Culham, J., Humphrey, G.K., Milner, A.D. and Goodale, M.A. (2003) 'Ventral Occipital Lesions Impair Object Recognition but not Object-Directed Grasping: an fMRI study', *Brain*, 126, 2463-2475
- Jeannerod, M., Decety, J. and Michel, F. (1994) 'Impairment of Grasping Movements Following a Bilateral Posterior Parietal Lesion', *Neuropsychologia*, 32, 369-380
- Jeannerod, M. and Jacob, P. (2005) 'Visual Cognition: A New Look at the Two Visual Systems Model', *Neuropsychologia*, 43, 301-312
- Kohler, I. (1964) 'Formation and Transformation of the Perceptual World', *Psychological Issues*, Vol.3 (4, Monograph number 12), 1-173
- Lamme (2006) 'Towards a True Neural Stance on Consciousness', *Trends in Cognitive Sciences*, 10(11), 494-501
- Laurence, S. & Margolis, E. (2003) 'Concepts and Conceptual Analysis,' *Philosophy & Phenomenological Research*, 67, 253-282.

- Lewis, D. (1995) 'Should a Materialist Believe in Qualia?', *Australasian Journal of Philosophy*, 73(1), 140-144
- Loar, B. (1990) 'Phenomenal states', *Philosophical Perspectives*, 4, 81-108
- Lycan, W. (1996) *Consciousness and Experience*, Cambridge, MIT Press
- MacDonald, G. (1992) 'The Nature of Naturalism', *Proceedings of the Aristotelian Society Supplement*, Volume 66, 225-244
- Matthen, M. (2005) *Seeing, Doing and Knowing*, NY, Oxford University Press
- McDowell, J. (1996) *Mind and World*, Cambridge, Harvard University Press
- _____ (1998) 'Having the World in View: Kant, Sellars and Intentionality', *Journal of Philosophy*, xcv, 431-491
- McGinn, C. (1991) 'Can We Solve the Mind-body Problem?', *Mind*, Volume 98, Issue 391, 349-366
- McGlinchey-Berroth, R. (1997) 'Visual Information Processing in Hemispatial Neglect', *Trends in Cognitive Sciences*, 1(3), 91-97
- McLaughlin, B. (2003) 'Color, Experience and Color Experience' in Smith, Q. (ed.) *New Essays on Consciousness*, Oxford, Oxford University Press.
- Milner, A.D. and Goodale, M. (1995) *The Visual Brain in Action*, Oxford, Oxford University Press
- Mozaz-Garde, M. and Cowey, A. (2000), 'Deaf Hearing: Unacknowledged Detection of Auditory Stimuli in a Patient with Cerebral Deafness', *Cortex*, 36, 71-80
- Myin, E. & O'Regan, J.K. (In Press): Situated perception and sensation in vision and other modalities: a sensorimotor approach. In P. Robbins and M. Aydede (Eds.), *Cambridge Handbook of Situated Cognition*. Cambridge, UK: Cambridge University Press.
- Noë, A. (2002) 'On What We See', *Pacific Philosophical Quarterly*, 83(1), 57-80
- _____ (2004) *Action in Perception*, Cambridge, MIT Press
- _____ (2005) 'Real Presence', *Philosophical Topics*, 33(1), 235-264
- _____ (2007) 'Magic Realism and the Limits of Intelligibility', *Philosophical Perspectives*, 21(1), 457-474
- Noë, A. and O'Regan, J.K. (2001) 'A Sensorimotor Theory of Vision and Visual Consciousness', *Behavioural and Brain Sciences*, 24, 939-1031
- Nagel, T. (1974) 'What is it Like to be a Bat?' *Philosophical Review*, 83, 435-450
- O'Regan, J.K. and Myin, E. (2002) 'Perceptual Consciousness, Access to Modality and Skill Theories', *Journal of Consciousness Studies*, 9(1), 27-46
- Palmer, S. (1999) 'Colour, Consciousness and the Isomorphism Constraint', *Behavioural and Brain Sciences*, 22, 923-989

- Palmeri, T.J, Blake, R, Marois, R, Flanery, M.A. and Whetsell, W. (2002) 'The perceptual reality of synaesthetic colours', *Proceedings of the National Academy of Sciences*, 99, 4127-4131
- Papineau, D. (2006), 'Phenomenal and Perceptual Concepts', in Alter, T. and Walter, S. (eds) *Phenomenal Concepts and Phenomenal Knowledge*, Cambridge, MIT Press
- Peacocke, C. (1992) *A Study of Concepts*, Cambridge, MIT Press
- _____ (2001) 'Does Perception have a Nonconceptual Content?' *Journal of Philosophy*, 98(5), 239-264
- Perry, J. (2001) *Knowledge, Possibility and Consciousness*, Cambridge, MIT Press
- Pettit, P. (2003) 'Looks as Powers', *Philosophical Issues*, 13(1), 221-252
- Philipona, D. and O'Regan, J.K. (2006) 'Colour Naming, Unique Hues and Hue Cancellation Predicted from Singularities in Reflectance Properties', *Visual Neuroscience*, 23 (3-4), 331-339.
- Poston, T. (2007) 'Cognitive Abilities and the Conceptualist/Nonconceptualist Debate', APA Pacific Division Meeting, April 3-8, 2007
- Prinz, J. (2000) 'A Neurofunctional Theory of Visual Consciousness', *Consciousness and Cognition*, 9(2), 243-259
- Putnam, H. (1960) 'Minds and Machines', in Putnam (1975) *Mind, Language and Reality*, Cambridge, Cambridge University Press
- Quine (1951) 'Two Dogmas of Empiricism', reprinted in Stainton, R. (ed) (2000) *Perspectives in the Philosophy of Language*, Canada, Broadview Press
- Radcliffe, M. (2008) 'The Phenomenology and Neurobiology of Moods and Emotions', in Gallagher, S. and Schmicking, D. (eds) *The Handbook of Phenomenology and Cognitive Science*, Springer
- Ramachandran, V.R. and Hubbard, E.M. (2001) 'Psychophysical investigations into the neural basis of synaesthesia', *Proceedings of the Royal Academy of London*, 268, 979-983
- _____ (2003) 'Hearing Colours and Tasting Shapes', *Scientific American*,
- Rauschecker, J.P and Tian, B (2000) 'Mechanisms and Streams for Processing of 'What' and 'Where' in Auditory Cortex', *Proc. Natl. Acad. Sci. USA*, 97, 11800-11806
- Rich, A.N. and Mattingley, J.B. (2002) 'Anomalous Perception in Synaesthesia: A Cognitive Neuroscience Perspective', *Nature Reviews Neuroscience*, 3, 45-52
- Rorty, R. (1965) 'Mind-Body Identity, Privacy, and Categories', *Review of Metaphysics*, 19, 24-54
- _____ (1979) *Philosophy and the Mirror of Nature*, Princeton University Press
- Rossetti, Y. (forthcoming) *Is Optic Ataxia a Visuo-motor Disorder? Perception, Action in the Visual Dorsal Stream*, forthcoming in the proceedings of the 'Perception, Action and Consciousness: Sensorimotor Dynamics and Dual Vision' conference, University of Bristol 2007

- Rossetti, Y, Rode, G. and Boisson, D. (1995) 'Implicit Processing of Somaesthetic Information: A Dissociation Between Where and How?' *Neuroreport*, 6, 506-510
- Sagiv, N. and Robertson, L.C. (2005) 'Synaesthesia and the binding problem' In Robertson, L.C. and Sagiv, N. (Eds), *Synaesthesia: Perspectives from Cognitive Neuroscience*. New York: Oxford University Press
- Schellenberg, S. (2007) 'Action and Self-location in Perception', *Mind*, 116, 603-632
- Schenck, T. and Milner, A.D, (2006) 'Concurrent Visuomotor Behaviour Improves Form Discrimination in a Patient with Visual Form Agnosia', *European Journal of Neuroscience*, 24(5), 1495-1503
- Schneider, G.E. (1969) 'Two Visual Systems', *Science*, 163, 895-902
- Schwitzgebel, E. (2002) 'Why did we Think we Dreamed in Black and White?', *Studies in History and Philosophy of Science*, 33, 649-660
- Schwitzgebel, E. (2008) 'The Unreliability of Naïve Introspection', *Philosophical Review*, 117(2), 245-273
- Sellars, W. (1997) *Empiricism and the Philosophy of Mind*, Cambridge, Harvard University Press
- Shoemaker, S. (1982) 'The Inverted Spectrum', *Journal of Philosophy*, 7, 357-381
- Smilek, D. and Dixon, M.J. (2002) 'Towards a synergistic understanding of synaesthesia', *PSYCHE*, 8(01)
- Taira, M., Mine, S., Georgopoulos, A.P., Murata, A. and Sakata, H. (1990) 'Parietal Cortex Neurons of the Monkey Related to the Visual Guidance of Hand Movements', *Experimental Brain Research*, 83, 29-36
- Thompson, E. (2007), *Mind in Life*, Harvard, Harvard University Press
- Toribio, J. (ms.) 'Is Perception Compositional?'
- Tye, M. (1995) 'Ten Problems of Consciousness', Cambridge, MIT Press
- _____ (2001) 'Consciousness, Colour and Content', Cambridge, MIT Press
- _____ (2003) 'A theory of phenomenal concepts', in O'Hear, A. (ed) *Minds and Persons*, Cambridge University Press
- Ungerleider, L. and Mishkin, M. (1982) 'Two Cortical Visual Systems', In Ingle, D.J., Goodale, M. and Mansfield, R.J.W. (eds.) *Analysis of Visual Behaviour*, Cambridge, MIT Press
- Varela, Thompson and Rosch (1991) *The Embodied Mind*, Cambridge, MIT Press
- Wallhagen, M. (2007) 'Consciousness and Action: Does Cognitive Science Support (Mild) Epiphenomenalism?', *British Journal for the Philosophy of Science*, 58(3), p.539-561
- Weiskrantz, L. (1998) *Consciousness Lost and Found*, Oxford, Oxford University Press