

Naturalising Phenomenology

Chapter 7

Introduction

So far I have shown that naturalists could, *in principle* accept the existence of subjective facts, but it remains to be shown *how* they can do so. This chapter will take up the remaining challenge. I have set out two lines of arguments in defence of the claim that there are subjective facts. First it was argued that we must appeal to subjective facts to account for the contents of experience. A large part of the chapter will be taken up with developing a naturalistic account of the contents of experience. The second argument for subjective facts – the argument from the experience of content – was given in the course of returning an answer to the question of why there should be something rather than nothing it is like to be a conscious creature. I argued that all conscious experiences are also pre-reflectively self-conscious. Thus the second goal of this chapter will be to sketch a naturalistic account of pre-reflective self-consciousness.

In the first section I shall explain how I understand the project of naturalising phenomenology. Some naturalistically inclined phenomenologists have proposed the use of phenomenology in the design of experiments. Subjects are instructed in the methods of phenomenology which they then put to use to arrive at first-person data. This first-person data is then employed to interpret data from brain scans recorded during the experiments.¹

While this strikes me as a fascinating and potentially fertile approach to the study of consciousness, what I have in mind by naturalised phenomenology is something different. I have argued that a naturalistic theory of mind must make room for subjective facts, if it is to successfully close the explanatory gap. I have given two reasons for believing in the existence of subjective

¹ For a recent discussion of this kind of approach see Gallagher (2003b). The inspiration behind this methodology comes from Varela's idea of a Neurophenomenology, see Varela (1996) and the lengthy introduction to Petitot *et al* (1999) for a description of the general approach.

Naturalising Phenomenology

facts. Naturalised phenomenology as I shall understand it has as its goal the assimilation of subjective facts.

In section 2 I return to the account of intentionality I have given in chapter 4. There it was argued that there is a distinctively phenomenological species of intentionality. A similar notion of intentionality has recently received discussion under the heading of “phenomenal intentionality”. Phenomenal intentionality is taken to be a species of narrow content which supervenes on a subject’s neurophysiological states. In section 3 I will argue that some caution is required in treating phenomenal intentionality as a variety of narrow content. Sections 4 and 5 examine whether the enactive view of perception can supply an account of phenomenal intentionality.² I argue that it can but only supplemented with a richer account of how we can perceive what J.J. Gibson called “affordances”. I will also argue that such an account can remain neutral on the question of whether phenomenal intentionality is a species of narrow content.

Central to what I called the argument from the contents of experience is the thesis that intentionality originates in our being-in-the-world. If the theory of intentionality I have advanced is to be shown to be consistent with naturalism, I must show how the notion of being-in-the-world can be assimilated within a naturalistic theory of mind. In particular I must make sense of the claim that our understanding of being can literally shape or constitute the objects of our experience. In section 6 I argue that this conception of experience calls for a radical overhaul in how think about the function of perception. We can no longer think of perception as serving the function of *detection*. Instead we must instead understand perception as serving a function for an animal in the context of a particular environment or

² The enactive approach to the study of visual perception has as its origins in J.J. Gibson’s ecological theory of perception. The term was first used in Varela et al (1991) and has recently found a powerful defence in the work of O’Regan and Noë, see for instance their (2001) and Noë (2004).

Naturalising Phenomenology

milieu. This understanding of perception is, I shall argue, in conflict with any view which takes perception to be detection. However such a revision in our understanding of perception will clear the way for a naturalistic account of being-in-the-world.

In section 7 I take up the challenge to the naturalist to supply an account of pre-reflective self-consciousness. In the final part of this chapter I return to the idea that the subject is implicitly represented in experience. I suggest that recent work on the role of motor imagery in skilful behaviour may provide us with the beginning of a story to tell about what it is for a subject to be implicitly represented in experience.³

By the end of this chapter I will have introduced empirical research which indicates that much of what the phenomenologist claims about the nature of conscious experience can be assimilated by naturalism. If the argument of this chapter is correct, there is a view of the mind already available to the naturalist which has found room for subjective facts. I have argued that making room for subjective facts is the key to closing the explanatory gap. It follows that cognitive scientists and neuroscientists have already begun to discover the tools for closing the explanatory gap. To the extent that it has seemed otherwise to some philosophers, this is perhaps because they have failed to attend to the phenomenology of conscious experience closely enough. They have brought to bear prejudices which are not borne out by close attention to phenomenology.

1. The Project of Naturalising Phenomenology

Let us begin by returning to the idea of a naturalised phenomenology introduced in passing in my introduction. Francisco Varela (1996) sets out the following working hypothesis for the project of naturalising phenomenology.

³ However we shall see that there is some reason to doubt whether this can be the whole story, for there is reason to think that a conscious experience could be PRSC even in the absence of motor imagery.

Naturalising Phenomenology

According to Varela phenomenology should be understood as standing in a relation of “reciprocal constraint” to the sciences of the mind.⁴ Phenomenology supplies detailed descriptions of conscious perceptual experience and the ways in which it is given form by our being-in-the-world. A naturalised phenomenology, according to Varela, will use these descriptions as data for uncovering new third-person descriptions of the physiological basis for consciousness.⁵ Phenomenology constrains the science of consciousness by supplying it with descriptions of conscious experience which it is the task of science to explain. At the same time the descriptions phenomenology supplies of experience, are also constrained by what the science of the mind has to tell us about the nature of perception.

Varela characterises phenomenology as providing a disciplined method for studying consciousness from a first-person perspective. Varela proposes that scientists train subjects in the use of the phenomenological reduction. Subjects can then use the phenomenological reduction as a window onto their own minds, in order to obtain first-person data. Varela conceives of a science of consciousness which uses this data in conjunction with third-person data from neuroscience to construct models of conscious experience. The distinctive contribution the practice of phenomenology can make is to contribute data to the science of mind that might otherwise go unnoticed.

This general methodology has recently been put into practice by the neuroscientist Antoine Lutz.⁶ Lutz and his colleagues exposed subjects to a 3D perceptual illusion. Subjects were asked to fixate for several seconds on a dot pattern containing no depth cues. At the end of this period the pattern was changed to one with binocular disparities. Subjects were asked to press a button as soon as they perceived the emergence of a 3D shape. EEG

⁴ See Varela (1996: 351)

⁵ See Gallagher (2003b); Lutz and Thompson (2003) and Borrett, *et al* (2000) for accounts of naturalised phenomenology along these lines.

⁶ For a discussion of these experiments see Lutz *et al* (2002), and Lutz and Thompson (2003)

Naturalising Phenomenology

signals were recorded throughout the trial. Immediately after subjects had pushed the button, they were asked to describe their experience. In these reports subjects would use what Lutz calls “phenomenal categories” which the subjects had devised for describing their experiences during a prior training session. During this training session, subjects were asked to direct their attention to their own mental processes during the task and to the ‘felt-quality’ that accompanied the emergence of the 3D image (Lutz and Thompson, 2003: 43). Lutz characterises this redirection of attention to experience as it is lived through, as corresponding to the phenomenological reduction. Thus he thinks of the phenomenal categories the subjects were later to make use of in their reports, as corresponding with perceptual invariants discovered through the practice of something akin to the phenomenological reduction. ‘In dialogue with the experimenters’, he tells us ‘(subjects) define their own stable experiential categories of phenomenal invariants to describe the main elements of the subjective context in which they perceived the 3D shapes.’ (*op cit*, 44)

Following the experiments, the reports the subjects made using these phenomenal categories were used to interpret the EEG data. Lutz hypothesised that each phenomenal invariant would be ‘characterised by distinct dynamical neural signatures before stimulation’ reflecting the degree to which the subject was ready for the emergence of the 3D shape. These ‘signatures would’ it was hypothesised ‘condition the neural and behavioural responses to the stimulus’. (*op cit*, 45)

Lutz *et al* did indeed find interesting correlations between the phenomenal categories subjects employed in their reports and patterns in brain activity as measured using EEG before the stimulus. They found that the states of preparation in subjects ‘modulated...both the behavioural response and the dynamical neural responses after the stimulation’. In addition they found that the ‘shape of the synchrony patterns’ while varying across subjects ‘were

Naturalising Phenomenology

stable in individual subjects throughout several recording sessions'. (*op cit*, p.46)

While I certainly do not mean to diminish the importance of the work that Lutz and his colleagues are carrying out, it doesn't seem right to me to describe this work as phenomenological. This research is better thought of as establishing the importance for the scientific study of consciousness of data arrived at through the careful use of introspection.

As I understand Husserl, he employed the method of bracketing in order to describe the different kinds of intentionality characteristic of consciousness, and how these different kinds of consciousness constitute the world we experience and think about. Husserl uses the term "constitution" to capture the sense in which consciousness contains the conditions necessary and sufficient for the existence of the objects we experience.

Lutz's experiment certainly isn't studying anything akin to constitution as Husserl understands it. Perhaps Lutz was thinking of the emergence of the 3D figure as something akin to constitution. However the reason subjects come to perceive a 3D figure in this experiment is because of the binocular disparities present in the dot pattern they perceive. The subjects do not come to perceive a 3D figure through constituting acts of consciousness of the kind Husserl describes. Lutz may have devised a disciplined method for studying variations in subjective experience. He may have also discovered interesting correlations between these variations in subjective experience and brain activity. Striking as these findings are, they don't have very much to do with phenomenology as it was practiced by Husserl, and the existential phenomenologists.

I take phenomenology to have an important contribution to make to the scientific study of consciousness through the descriptions it gives of intentionality, and how our conscious experiences come to exhibit intentionality. It is these descriptions of intentionality which naturalists must

Naturalising Phenomenology

assimilate if they are to provide a satisfactory account of consciousness. The phenomenologist's account of intentionality entails the existence of subjective facts. It follows that naturalism can assimilate the account of consciousness we find in phenomenology only if it can assimilate subjective facts. It is this task that will occupy our attention across the next two sections.

2. The Phenomenal Intentionality Thesis

Naturalistic theories of intentionality have often supposed that 'if intentionality is real, it must really be something else'.⁷ They have supposed that it must be possible to give a reductionist account of intentionality which accounts for intentionality in non-intentional terms. Fodor (1987) captures the spirit of much work in this area when, in discussing what it will take to integrate intentionality into the natural order, he says: 'what we want at a minimum is something of the form '*R* represents *S*' is true iff *C* where the vocabulary in which *C* is couched contains neither intentional nor semantic expressions.' (*op cit*, p.32) I shall call accounts of intentionality which seek to identify non-intentional or non-semantic conditions necessary and sufficient for a state to count as a representation, "reductive" accounts of intentionality.

Reductive accounts of intentionality have thus far proven to be something of a failure. Their defects can be brought out by thinking of cases in which our thoughts or experiences fail to refer to anything. In cases of reference failure, there is nevertheless something that the thought or experience has the *goal* of representing correctly. What reductive accounts have so far failed to explain is what it is for a representation producing system to have a goal of this kind.⁸ No one has yet succeeded in identifying conditions which are necessary and sufficient for an organism to have the goal of representing correctly.

⁷ Fodor (1987: 97)

⁸ For a persuasive defence of this conclusion see Walsh (2002).

Naturalising Phenomenology

Must a naturalistic account of intentionality also be a reductive account of intentionality? I have claimed that naturalistic accounts come in two varieties; they either eliminate or assimilate some putative natural phenomenon. I am assuming that elimination isn't the appropriate strategy to adopt when looking for an account of intentionality. The appeal to intentional content does essential explanatory work in explaining our behaviour which cannot be done by explanations which do not attribute content-bearing states to an agent.⁹ That leaves naturalisation by assimilation. Must a naturalist seeking to assimilate intentionality aim for a reductive account of intentionality?

I will present a theory of intentionality which explains how our conscious experiences get their content by reference to certain skills or capacities of a perceiving animal. A reductive account of these skills and capacities would seek to characterise them in non-intentional terms. We will see that the skills and capacities in question must be understood as the skills of a particular animal acting in pursuit of certain goals. I doubt that there is any way of accounting for intentionality in naturalistic terms that doesn't somehow advert to a representational system and the goals it has, though this is not something I can attempt to argue for here. Still I think the naturalist can assimilate intentionality. The practical understanding I shall appeal to in explaining how an experience get its content attributes to an animal skills and capacities which are susceptible to explanation by cognitive scientists and neuroscientists. This practical understanding cannot however be understood apart from the goals and purposes of the organism which the animal acts to bring about. So even though we might not be able to give a wholly reductive definition of intentionality, we are nevertheless beginning to understand how

⁹ Many philosophers have attempted to argue for this conclusion, satisfactorily to my mind. See for instance Dretske (1988); Fodor (1989); Horgan and Woodward (1985); Rudder-Baker (1987 & 1995). To enter into this debate here would take me too far away from my concerns in this chapter which is to offer a naturalistic account of the conception of conscious experience I have proposed in previous chapters.

Naturalising Phenomenology

intentionality could be realised by conditions describable by the sciences of mind.

The question I will take up in what follows concerns what it is for a conscious perceptual experience to have an intentional content.¹⁰ I will not be proposing a more general account of intentionality. The intentional content that our conscious experiences carry is importantly different from beliefs, desires and other of the so-called propositional attitudes in that there is good reason to think that it is, at least partially, non-conceptual.¹¹ A child or animal, for instance, could very well perceive Clyde playing the piano or smell the toast burning (to borrow two examples from Dretske (1993/1997)) even if it did not possess the concept of a piano or of burning toast. More generally we can say that a conscious experience E has non-conceptual content if a creature can undergo E while not possessing the concepts required for specifying what it is that E represents. Supposing that experiences do have non-conceptual content, then the question arises of what it is for an experience to have a non-conceptual content. In particular what is it for a creature to represent P rather than Q when the creature lacks the concepts required for specifying what it is representing? This is one of the two questions that will occupy me for the remainder of this chapter.

The phenomenologist takes a perceptual experience to have an intentional content which is constitutively determined by its phenomenology. The first thing to note is that an experience's having a particular phenomenology need not place a requirement on the subject with respect to the possession of

¹⁰ I will be seeking to give a naturalistic account of what I called in chapter 2 "operative intentionality". I shall argue that perceptual content is operative intentional content. I will not attempt to account for what I called "cognitive intentionality", the kind of intentionality belonging to propositional attitudes.

¹¹ Whether an experience's representational properties could be entirely nonconceptual is an interesting question which I cannot take up here. Bermúdez (1995) draws upon work in developmental psychology to argue that an experience could have a content that is entirely nonconceptual. Peacocke (1994) disagrees, arguing instead that a representation has different layers of content some of which are nonconceptual and others of which are conceptual. Peacocke has however since changed his mind, see the appendix to his (1994/2003).

Naturalising Phenomenology

concepts. I take it that an infant and I share something in common when we smell the burning toast. We both undergo experiences with a phenomenology characteristic of smelling burning toast, though only I possess the concepts required for saying what I am experiencing.¹²

I don't mean to deny that concept-possession might be a necessary condition for enjoying *some* kinds of experiences. Think about the wine-taster who can discriminate qualities in the wine he is drinking that escape my palate. Arguably he is in possession of recognitional concepts that I haven't myself mastered, and this gives him a richer appreciation of wine than I myself enjoy. This is not to say that I don't enjoy wine, but it is to say that I don't enjoy wine in the same way as a wine-taster. He enjoys an experience with a richer phenomenology than me as a consequence of the recognitional concepts he has acquired.

An experience's phenomenology, I argued in chapter 4, is a joint product of (a) the entities and properties in the world which the experience seems to present to a subject, and (b) the ways in which an experience presents those entities and properties. (a) and (b) correspond to the contents of experience and mode of presentation respectively. Consider my experience of a round plate that looks elliptical to me. On the one hand my experience presents to

¹² There is a large question as to whether possession of concepts changes the character of a subject's experience across the board. McDowell (1994, lecture III) seems to want to claim that the fact that human beings are speakers of a language transforms the kinds of experiences we can enjoy as compared with other animals. (He adopts this thesis from Gadamer, see his lecture VI.) McDowell acknowledges that our experience shares something in common with that of non-human animals, which he characterises as 'perceptual sensitivity to our environment' (McDowell, 1994: 64). However he claims that we exhibit this perceptual sensitivity 'in a special form. Our perceptual sensitivity to our environment is taken up into the ambit of the faculty of spontaneity, which is what distinguishes us from them.' (op cit, 64) This leads McDowell to deny that there is a common component to the experiences we enjoy and those of non-human animals. I would like to agree with McDowell that the possession of concepts and a language *may* transform our experiences without denying that there is a layer of content to our experiences which is non-conceptual. I am also sympathetic to the claim that the kinds of experience enjoyed by animals may be very different from our own, though I am less inclined than he to put this difference down to the role of language and concept-possession. Rather the difference between the experiences of animals and our own has more to do with what non-human animals are interested in representing, for more on this point see my discussion in section 6.

Naturalising Phenomenology

me a plate that is shaped round. This is not however how the plate is represented in my experience. The way in which the plate is represented is such that it looks elliptical to me. I have suggested that my experience of something that is round but that looks elliptical is the joint product of the factors described in (a) and (b). My experience presents me with a round plate. However the way in which it does so is by presenting me with something that looks elliptical.¹³

An experience can of course have the phenomenology it does without successfully representing anything in the subject's environment. My experience can present to me something that seems to be a cat say when I am in fact experiencing a fox. I can even see something that seems to me to be a cat when there are no animals of any kind around whatsoever. Nevertheless in both these cases my experience *purports* to represent a cat. My experience purports to represent a cat insofar as it has certain representational properties – the properties characteristic of experiences that represent cats – in virtue of which it is assessable for truth or falsity.

Now consider again the claim I have attributed to phenomenologists which says that an experience's intentional content is constitutively determined by its phenomenology. This is to say that an experience is assessable for truth or falsity in virtue of what it is like to undergo this experience, where what it is like to undergo an experience is analysed in terms of the two conditions I have just described. I shall label this proposal "The Phenomenal Intentionality Thesis".

All theorists are agreed that sensory experiences have intentional features in virtue of which they are assessable for truth or falsity. What differentiates

¹³ My claim that there is a way in which an experience presents a plate say, when it looks elliptical doesn't require me to introduce any non-representational properties. A plate looks elliptical, I shall claim because of factors to do with the context in which it is being perceived. In this case it is because the plate is being viewed from a particular angle that it looks elliptical. The fact that it is being viewed from this angle forms a part of the viewing conditions. Here I am influenced by the account of perceptual constancy phenomena given by Kelly (2001/2003).

Naturalising Phenomenology

the phenomenal intentionality thesis from other theories of intentionality is the claim that a sensory experience is assessable for truth or falsity just in virtue of what is like for a subject to undergo this experience. Most naturalistic theories of intentionality suppose that there is some relation to the world which accounts for an experience having the intentional features it does. This relation might be causal covariation under ideal conditions, or a relation to the world that holds when the consumer systems that make use of representations to generate actions are functioning properly.¹⁴ The phenomenal intentionality thesis denies that a subject must stand in any particular relation to the world in order to be assessable for truth or falsity. All that is required for an experience to have the intentional features it does is for there to be something it is like for the subject to undergo this experience. What makes it the case that an experience has the particular correctness conditions it does is fully determined by the ways in which this experience presents the world as seeming to its subject.

Perhaps an example will help. Consider my experience as of a white rabbit. My experience presents me with a creature that is furry, has floppy ears and bucked teeth, is coloured white *etc.* These descriptions capture some of the ways in which my experience presents the rabbit as appearing. This experience will be accurate if there is indeed a rabbit present in my environment having the features I have just described. What is it for my experiences to have these intentional features? I claim my experience has these intentional features in virtue of the ways in which it presents the world as *seeming* to me. The world *seems* to me to contain a rabbit, and this is what makes it the case that I am undergoing an experience which *represents* the presence of a rabbit.

¹⁴ For the first kind of view see Stampe (1977) and for proponents of the second kind of view see Millikan (1984) and Papineau (1993).

Naturalising Phenomenology

In chapter 4 I distinguished between a strong and a weak reading of the phenomenal intentionality thesis. According to the strong reading, there is no relation to the world which our experiences depend on for their content. On this view an experience can present the world as seeming to be a certain way whether or not the world of material things exists.¹⁵ The weaker reading denies that there exists a particular relation to the world which a subject must stand in if s/he is to represent an object x. This is consistent with saying that a subject must stand in some relation to the world if her experience is to have the content it does.

In what follows I will be developing a naturalistic account of perceptual intentionality which endorses the weaker reading of the phenomenal intentionality thesis. I will propose an account of perceptual intentionality which takes the contents of perception to be constitutively determined by phenomenology while at the same time allowing that an experience's phenomenology may depend on factors outside the subject's head. In the next section I will say something about why I think a naturalist seeking an account of phenomenal intentionality ought to favour the weaker reading of the phenomenal intentionality thesis over the stronger reading. My aim will be to bring out exactly in what sense an experience's phenomenology could involve things in the subject's local environment.

Before I consider the strong reading in detail it is worth noting that if it were correct, the project of assimilating phenomenal intentionality into a naturalistic account of the mind would turn out to be very difficult undertaking indeed. Not only would we have to explain how the brain generates phenomenal consciousness – Chalmers' hard problem. We would also have to explain

¹⁵ Loar (2003) and Horgan and Tienson (2002) both endorse the strong reading of the phenomenal intentionality thesis. In chapter 4 we saw that some caution is required in attributing the strong reading to Husserl even though he did on occasion come close to defending such a position. For an interpretation of Husserl which does attribute the strong reading to him see for instance Dreyfus (1982) and McIntyre (1982).

Naturalising Phenomenology

how the brain could generate states that exhibit intentionality intrinsically.¹⁶ I shall present an account of phenomenal intentionality which shows how it might be possible for the brain to achieve this. However the biological plausibility of this account increases considerably once we allow that the brain might not be doing all of this work on its own. I will argue that an account of phenomenal intentionality ought to allow room for the environment and a creature's embodiment to also play a role in generating its intentional states.

3. Against the Strong Reading of the Phenomenal Intentionality Thesis

It is natural to think that an experience has the phenomenal character it does regardless of whether the experience is veridical or not. The infamous argument from hallucination rests on the possibility of a subject undergoing a hallucinatory experience which is qualitatively indistinguishable from a veridical experience.¹⁷ To say that two experiences are "qualitatively indistinguishable" is to say that a subject undergoing the two experiences would be unable to detect any difference between them. To illustrate consider a subject enjoying a veridical experience of an apple. We seem to be able to conceive of this subject's brain being tampered with so that when he blinks the apple could be destroyed without the subject noticing any change in his experience. In order for our subject to continue to undergo an experience as of an apple all that seems to be necessary is for us to hold constant whatever neurophysiological activity realises his experience. Thus it would seem that a subject could undergo an experience with the phenomenology characteristic of a veridical experience of an apple without standing in a relation to any existing apple. If what it is like to undergo an experience does suffice to

¹⁶ Horgan and Tienson (2002) end their paper by noting the following implication of their view: 'conscious intentional states are intrinsically, by their very nature, directed toward whatever they are directed toward. Thus, the hard problem includes this: why should a mental state that is grounded in this physical or physical/functional state be by its intrinsic phenomenal nature directed in this precise manner? And this is a very hard problem indeed.' (Horgan and Tienson 2002: 530)

¹⁷ See Valberg (1992) for an excellent discussion of the argument from hallucination.

Naturalising Phenomenology

determine an experience's intentional features, it would seem that an experience can have the intentional features it does independent of any relation a subject stands in to the world.¹⁸

This argument relies on the questionable assumption that there are neural correlates of conscious experience with contents that, in a sense to be explained, 'match' the contents of experience.¹⁹ Noë and Thompson (2004) define a neural correlate of consciousness as a minimal neural representational system N such that (1) N is sufficient for the occurrence of some experience E and (2) there is a match between the content of N and E.²⁰ There are really two claims that the above argument assumes, both of which can be challenged. The first assumption is that there is a system of neural activity, N which is *sufficient* for the subject to undergo his experience as of an apple. I'll call this "the neural substrate assumption". The second assumption is that this neural substrate has a representational content which matches the content of the conscious experience a subject enjoys. I'll call this the "matching content assumption".

The neural substrate assumption implies that a subject's embeddedness in a particular kind of environment plays no role in determining an experience's phenomenology. A subject could be a brain in a vat and enjoy experiences with just the same phenomenology. Many philosophers have believed that something along these lines is true, but whether they were right would seem to be an empirical question. One empirical possibility that surely shouldn't be ruled out from the armchair is that the representational vehicles underlying our

¹⁸ This is a conceivability argument and like any argument of this kind it is subject to the normal kinds of worries one might have about the relation between conceivability and possibility. Wilson (2004: ch.10) exploits a point of this kind in arguing against the phenomenal intentionality thesis. He attacks the notion of phenomenal intentionality because he takes it to entail individualism – the view that intentional content is to be individuated narrowly. I shall be defending a notion of phenomenal intentionality which is compatible with Wilson's attack on individualism. Thus I can fully concur with his attack on Horgan and Tienson's strong reading of the phenomenal intentionality thesis without giving up on the idea of phenomenal intentionality altogether.

¹⁹ Here I will draw on arguments to be found in Noë & Thompson (2004).

²⁰ They attribute this definition to Chalmers (2000) but I have not been able to read this paper.

Naturalising Phenomenology

experience extend into the world. This is to say that the conscious experiences we enjoy could be the result of dense and continuous interactions taking place between body, brain and an animal's surrounding environment.²¹ On this understanding of conscious experience, the environment in which an animal is embedded will make a significant contribution to the character of a subject's experience. It will form a part of the feedback loop which on this view is the minimal sufficient condition for a subject's enjoying a conscious experience.

The moral is that only further empirical research will enable us to establish whether there is a neural substrate which suffices to give us the kinds of experiences we enjoy. It is empirically possible that there is no neural substrate sufficient to generate conscious experience, but that instead our experiences are identical with representational vehicles that extend into the world.

One more comment on the superficial appearance of conceivability that attaches to hallucination scenarios. Whatever plausibility attaches to these thought experiments may well derive from the fact that something else, a demon or some other kind of deity, a super-intelligent scientist, a supercomputer, is enlisted to construct an environment which resembles our own. It is an empirical question whether it is our brain which ordinarily constructs this appearance for us or whether the environmental setting in which a creature is embedded itself plays an ineliminable role in giving us the kinds of experiences we enjoy. If the latter, then brain-in-the-vat stories are only conceivable to the extent that some other virtual environment can be created to play the role currently played the environments we inhabit.²²

²¹ Hurley (1998: ch.6); Noë (2004, ch.7 & 2005); Rowlands (2002 & 2003: ch.10); Thompson and Varela (2001); and Wilson (2004, ch.9) all defend an externalist account of the vehicles of conscious experience, "externalist" in that it takes the vehicles in question to extend into the world, through feedback loops that essentially involve interactions with things in the subject's external environment.

²² Noë (2004, ch.7: 218) makes what I think is the same point.

Naturalising Phenomenology

This brings me to the matching content assumption. The matching content assumption says that for any given experience E there is a system of neural activity N with a representational content that matches E's representational content. What does it mean to say that E and N match in content? Returning to the subject looking at the apple, the idea is just that by keeping the activity in neural system N constant, we could continue to cause in this subject an experience as of an apple. The experience as of an apple we manage to reproduce is, we are to suppose, qualitatively indistinguishable from a veridical experience of an apple. Thus we can say that N and E match in content if whatever features E represents, N also represents. Given that we can produce an experience E just by holding the activity constant in N, it is natural to think that a match in content of this kind does indeed obtain.

However again there are empirical reasons for being suspicious of such a claim. The matching content assumption seems to commit us to the existence of a system of neural activity which is the bearer of a representational content as rich as the contents of experience. The first point to note is that there is some controversy about how rich the contents of experience are. When Ernst Mach (1886) drew a picture of himself stretched out in his armchair, the picture he drew was of a visual field, uniform and high resolution in its detail. However there are number of reasons for doubting whether this accurately reflects what our visual fields are like.²³ Consider our visual perception of colour, for instance. We hardly perceive the objects at the margins of our visual field to be coloured at all, whereas the objects that are at the centre of visual field appear to be coloured in splendid detail. Alva Noë tells us:

'Rods and cones are not evenly distributed across the surface of the retina. Outside the high resolution central (foveal) region, there are increasingly few cones. As a result of this, the eye is nearly colour-blind in

²³ Here I am indebted to arguments in Noë (2004, ch.2).

Naturalising Phenomenology

its parafoveal region. Despite these defects we do not experience the world as black and white at the edges.’ (Noë, 2004: 37)

Consider for example my experience of the colour of the carpet I am standing on which is a uniform muddy brown. Do I experience the muddy brownness of the carpet at the periphery of my visual field just as much as I experience its muddy brown at the point in my visual field which I am currently fixating? Suppose we say I do. In order to explain how this could be we have to hypothesise mechanisms in the brain which correct for the limitations in the image of a scene that falls on our retina at any moment. It is the matching assumption that commits us to a hypothesis of this kind. If we suppose that what we experience is a scene uniformly detailed in the information it gives us about the colours of things, it is tempting to think this is because there is a system of neural states with a content that matches that of our experience. It is this system of neural states which gives me an impression of the carpet I am currently standing on as being uniformly muddy brown all the way around me.

There is however another possibility. Perhaps our experience doesn't represent a world that is high resolution and uniform in its detail from the centre of our visual fields to their margins. All the detail is there in the world ready to be tapped as when we need it, simply by the movement of our head or body, so why do we need to construct detailed representations of the way things are in the world at each moment? Surely we need to represent only as much detail as is required for what we are doing. What we visually perceive at any given moment may be function of what we have our gaze fixed upon at the time.²⁴

²⁴ O'Regan (1992) has suggested that a creature's local environment may be thought of as functioning as kind of external memory which is available to be made use of by the senses, just as (some of) our memories are available for recall, as when we need them. This idea is also discussed in O'Regan and Noë (2001: 946).

Naturalising Phenomenology

That something along these lines might be true is borne out by the phenomena of change-blindness. Usually when a change takes place in a visual scene our attention is alerted to the change by some kind of flicker of movement. This led to the prediction that if we could somehow be prevented from noticing the flicker of movement we could also fail to notice changes, even when they were happening right in front of our eyes.²⁵

In one striking study Simons and Levin (1998) showed that subjects could fail to detect a change in a person they are giving directions to. Hayhoe, Bensinger & Ballard (1998) asked subjects to perform a copying task involving blocks. In this experiment subjects regularly failed to notice a change in the model they are copying. Perhaps the most well-known study of this kind (Simons and Chabris 1999) subjects were given the task of watching a basketball game and watching the number of times one or the other team takes possession of the ball. Many subjects performing this task failed to notice the presence of a man in a gorilla suit strolling through the centre of play.

The most conservative reading we can give of these studies suggests that if our brains are in the business of producing detailed representations of our environments from moment to moment, the information these representations carry may not be available for the formation of memories and the making of verbal reports. To the extent that we think we have a conscious experience of a scene in all its glorious detail, this conservative reading would have it that we are mistaken. What we consciously perceive in a visual scene at any given moment will be as much or as little as we are attending to in the visual scene at that moment.

A more radical reading of these experiments takes them to count as evidence against the brain constructing rich internal representations of our

²⁵ For a discussion of these experiments and their implications see Noë, Pessoa, and Thompson (2000).

Naturalising Phenomenology

environment from moment to moment. On this reading, we fail to detect the changes in these cases because those changes aren't represented by us. I am not going to attempt to decide between these two ways of thinking about the change blindness results. Either way change-blindness seems to have significant implications for the matching assumption. The change-blindness experiments seem to conflict with the following claim:

(1) Conscious experiences seem to present us with a world in all of its detail, and this detail is reproduced by the representations our brains construct of a visual scene.

This is just a statement of the matching assumption. Thus at first glance the change-blindness would seem to come into conflict with the matching assumption. There are at least two ways of responding to the falsity of (1).²⁶

We could say:

(2) We are mistaken about how our experiences seem to us (see Dennett (2002)). There is no detail in our conscious experience, we just think there is. What the change blindness experiments show is that we don't see that world in all its detail. They highlight just how wrong we can be about our conscious experiences. It seems to me that this interpretation of change-blindness is quite consistent with the matching assumption. One could say that there is a neural representational system N the content of which matches the contents of visual experience. We might say that mistakes arise when we come to say or judge exactly what it is we are seeing, and it is this mistake that the change-blindness experiments have brought to light.

²⁶ Actually there is a third possibility I can think of. We could say that neither our conscious experiences nor our neural representations are detailed in the information they give us about a visual scene. The change blindness experiments show that we are mistaken about the character of our experiences. They show this by showing that there are changes that take place which our brains do not represent. If this is right our neural representations are as impoverished in the information they carry as our conscious experiences. Thus the matching assumption holds. The difficulty this response faces is that it doesn't explain the mistake it attributes to us, whereby we take our experiences to present a visual scene to us in all of its detail. The advantage of (2) and (3) is that they at least attempt to do justice to the ways in which our experiences ordinarily present the world.

Naturalising Phenomenology

There is another possible response however, which isn't consistent with the matching assumption. It says:

(3) There is as much detail in our conscious experiences as there has always seemed to be but there is very little in the way of detailed information in the representations our brains produce. Our experiences seem to present us with a world rich in information, and this is exactly what they do. All of the detail we seem to experience is out there in the world. Our sense that we experience this detail is a reflection of the fact that this detail is there for us to explore using our senses. Thus, the matching assumption is false. Our experiences are richer in content than the neural representations which enable those experiences.

In what follows I will assume that the sense we have of seeing the world in all its detail isn't an illusion. I am going to assume that this is a datum that needs to be explained. What the change blindness experiments show is that we don't perceive all of the information there is in the world in one go.²⁷ The results of these experiments provide evidence against the claim that the brain is constantly engaged in constructing what we might call a snapshot of a scene. While a snapshot would capture all of the detail in a scene including any changes that might take place from moment to moment, the representations our brain produces do not seem to be like this. The possibility I want to take up next is that the brain doesn't generate phenomenal intentionality on its own. Rather our experiences get their phenomenal intentionality from a variety of practical understanding or knowledge that accompanies all of our conscious experience. This understanding might consist in part of knowledge of how movement will affect the ways in which an object is sensed by us. While it is undoubtedly true that the brain constitutes a necessary or enabling condition in this account of phenomenal intentionality I shall propose, the account I will give is consistent with saying that neural activity isn't a sufficient condition for phenomenal intentionality. Thus the

²⁷ Here I am following Noë, Pessoa and Thompson (2000) and Noë (2004: ch.2).

Naturalising Phenomenology

account I shall give of phenomenal intentionality is incompatible with the strong reading of the phenomenal intentionality thesis.

4. Husserl and Merleau-Ponty on Phenomenal Intentionality

In the last section we encountered a puzzle. Our experiences seem to present us with a world elaborate in its detail, and yet the visual representations our brain constructs do not seem to reproduce anything like all of this detail. What we need to explain then is why it seems to us that we experience all that detail given that it doesn't look like this is something made available by the internal representations our brain produces.

A good starting point for addressing this question is to return to a related puzzle which we saw Husserl raise in chapter 4. There we saw Husserl point out that an object is never fully given in experience: there is always more to an object than it presents to us at any given moment.

'In principle a physical thing can only be given "one-sidedly...A physical thing is necessarily given under simple modes of appearance...but what is "actually" predelineated is accompanied by a horizon, by a more or less vague zone of indeterminateness.' (Husserl, 1913/1972 §44, 100)

What Husserl has in mind is the uncontroversial fact that when I look at a house say, what I see of the house will depend on where I am standing in relation to it. Furthermore, I cannot see the entire house at once. All I see of the house at any given moment will be the sides of the house that are visible to me, and there will always be other aspects of the house that aren't visible to me at the time. I nevertheless have a sense of the whole house as present to me in my perception.²⁸ How can this be?

²⁸ If you are not persuaded of this, consider the difference in what it is like to see something you know to be a stage prop, a two-dimensional façade, and a genuine house. We can even suppose that the two-dimensional façade is an exact replica of the front of a house. Still there is a difference in what it is like to look at something one knows to be the façade of a house and what it is like to look at a genuine house. The difference resides not in what one is currently presented with in experience but in what one would expect to see if one were to

Naturalising Phenomenology

I take it that the puzzle Husserl has raised exactly parallels the question raised by the change blindness studies. These studies suggest that although we experience a world which seems elaborate in its detail, not all of this detail is reproduced in the visual representations our brains construct.²⁹ How is it that our experiences give us a sense of a world that goes beyond what we are representing at each moment?

The answer Husserl returns to this question doesn't seem to me quite right, but it will nevertheless be instructive to consider where he goes wrong. Husserl thought that the parts of the object that are not currently perceived are nevertheless represented in experience. Those aspects of the object which are not currently seen form what Husserl calls a 'horizon of indeterminateness' (see the above quote). In every perceptual experience Husserl tells us that 'the sides of the object which are actually perceived refer to sides which are not yet perceived but which are only anticipated...as aspects to come in perception.' (Husserl, 1931/1973: §19, 82)

Husserl's view seems to be that based on the side of the thing currently in view we form anticipations about the other sides of the thing that would come into view through our exploration of the thing. For Husserl the hidden aspects of things are not materially present in experience but they are nevertheless represented in the form of hypotheses about what I would see if I were to change my relation to the object. These hypotheses can be either confirmed or falsified by my further exploration of the thing. What is it for me to have an experience as of a lemon for instance when I am seeing only the lemon's facing side? Husserl would have us believe that I perform something akin to an act of interpretation whereby I take the part of the thing I am seeing to belong to a lemon. This interpretation I give can be confirmed or disconfirmed

move around the thing. In the case of a façade one expects to see something that is flat, whereas in the case of a genuine house one expects to find other sides of the house that were previously hidden from view come into one's line of sight.

²⁹ To be more precise, if all of this detail is reproduced in the internal representations the brain is busy producing, not all of this information is made available to consciousness.

Naturalising Phenomenology

by further exploration of the lemon. Perhaps it's a wax lemon, and I discover this by taking hold of it and feeling its weight, or by trying to cut into it.

Is it really the case that the hidden sides of a thing are represented in the form of expectations or hypotheses I have about what I would see if I were to change my spatial relation to a thing? Certainly if this were true it would explain why I have a sense of seeing a world in all its elaborate detail. I experience the world in this way as a result of the many expectations I am constantly forming about what my local environment contains.³⁰ Husserl's reply makes the object of perception into a sum of the actual and possible views one can take on it. The lemon as it is perceived by me is nothing but the actual parts of it I am currently presented with in my experience of it, and the other possible views which I hypothesise I could take on the lemon were I to alter the spatial relation I stand in to it. The intentional content of my experience – its purporting to represent a lemon – is in no way dependent on the existence of the lemon. I could form the hypothesis that the thing I am presented with is a lemon, and thereby undergo an experience that presents me with what seems to be lemon, no matter whether the lemon existed or not.

Of course this could be the way perceptual intentionality works – it could be that this is how perceiving creatures come to stand in an intentional relation to the objects of their experience. If so, it will turn out that phenomenal intentionality is indeed an intrinsic property of our conscious

³⁰ This is slightly misleading. Husserl is well aware that we cannot be constantly engaged in forming hypotheses about the absent aspects of a thing. He makes a distinction between what he calls "active" and "passive" synthesis by way of acknowledging this point. (See for instance the recently translated series of lectures from 1918-1926 devoted to this distinction, see Husserl (1970/2001). One of the last works Husserl wrote (1948/1973) also contains extensive discussions of this distinction.) Husserl describes the simplest forms of perception as involving what he calls "passive synthesis". Active synthesis by contrast is what a person engages in when he makes judgements about the objects of his experience – when he judges that the cover of the book is coloured red, say. Husserl claims that the process of forming hypotheses about a thing's hidden sides can also be passive, in that it needn't involve a person actively forming beliefs or making judgements about a thing's hidden properties. The expectations we form about what further exploration would reveal of a thing can be driven purely by our past and present encounters with a thing. We can, as a result of our past experiences of houses say, take the front side of the house we are looking at to indicate the presence of other sides of the house which are not currently visible to us.

Naturalising Phenomenology

experiences. We have a sense of a thing's presence including the parts of it that aren't currently in view, through the generation of hypotheses. Of course we have to take the side of a thing that is currently in view to 'refer' (in Husserl's words, see above) to other possible views one could take on the object. The hypotheses I form about the objects of my experience aren't completely unconstrained. However the representation that results is one that is formed independently of the relations I happen to stand in to the world.

Suppose we were to treat this as a serious empirical hypothesis about how we come to have experiences of the presence of three-dimensional material things given that what we start out with is a one-sided point of view on a thing. Understood in this light we could take Husserl to be claiming that our brains could achieve the representation of a full-fledged material thing independently of our relation to that thing or any other thing.³¹ This is something our brains could achieve through the formation of expectations or anticipations about what further exploration of the thing would reveal. If we were to take Husserl's descriptions of perception and turn them into an empirical hypothesis we could say that the brain is constantly engaged in producing something like a simulation of reality. This simulation is then tested out on the world through our actions.³²

Merleau-Ponty offers a different solution to our problem of how we can have a sense of being presented with a whole object when we are sensing only a part of it. Whereas Husserl takes the hidden sides of a thing to be perceptually absent, Merleau-Ponty takes them to be in an important sense

³¹ This is of course not what Husserl was saying; we saw in chapter 1 that Husserl was most definitely not a materialist.

³² For a recent proponent of such a view see Metzinger (2003: 50-62). He tells us that in his opinion: 'some of the best work in neuroscience' (here he cites Singer 2000 and Leopold and Logothetis 1999) 'suggests a view of the human brain as a system that *constantly* simulates possible realities, generates internal expectations and hypotheses in a top-down fashion, while being constrained in this activity by...a constant stimulus-correlated bottom-up stream of information, which then finally helps the system to select one of an almost infinitely large number of internal possibilities and turns it into a phenomenal reality.' (Metzinger, 2003: 51) This seems to me more or less what Husserl would say if he took the brain to be responsible for constituting our experiences of the world rather than the transcendental ego.

Naturalising Phenomenology

positively present in experience. Sean Kelly (2004) characterises the difference well:

'Husserl thinks that it is indeterminate, from the point of view of the current visual experience, what the features of the back-side of the object are. Merleau-Ponty by contrast thinks that my current visual experience contains something that is itself an indeterminate presentation of the back. For Husserl, it is not yet determined what I see; for Merleau-Ponty what I see is indeterminate.' (Kelly 2004: 81)

Merleau-Ponty claims that the hidden sides of the object form a part of what my experience represents, not only in the form of anticipations or expectations, as we have seen Husserl claim. A thing's hidden sides affect the movements I would direct towards it in reaching and grasping for that thing, for instance. In virtue of the role the hidden sides play in guiding my movements, Merleau-Ponty thinks we should say that the thing's hidden sides are presented in my experience *now*. They are not represented in the form of expectations I have about what my *future* exploration of thing would uncover. Rather they form a part of what I experience in the here and now. Merleau-Ponty talks of 'an indeterminate vision, a vision of something or other'³³ which 'is not without some element of visual presence.' (Merleau-Ponty, 1962: 6)

That the hidden sides of a thing are indeterminately present in experience is something that Merleau-Ponty thinks is borne out by the fact that the movements I direct towards the thing are sensitive to the hidden parts of the thing. How can the parts of an object which are currently not visible to me nevertheless make a contribution to how I direct my movements towards a thing? Merleau-Ponty's answer to this question appeals to a bodily

³³ Kelly (2004: 80-1) argues that the phrase '*vision de je ne sais quoi*' which is translated by Colin Smith as 'vision of something or other' ought to instead read 'vision of I do not know what'. He argues for reasons that will become clear in a moment that this better captures Merleau-Ponty's claim that 'we must recognise the indeterminate' – the hidden sides of a thing, say – 'as a positive phenomenon' (Merleau-Ponty, 1962: 6)

Naturalising Phenomenology

understanding I have of the thing. There is a sense in which I understand a thing's shape, size and weight through the knowledge I have of how to direct my movements towards that thing. This knowledge doesn't consist in my having reliably formed beliefs about exactly which movements would be appropriate to the object of my perception. In particular it doesn't involve my believing that the thing has certain hidden properties that are not currently in view, and that these hidden properties make certain behaviours appropriate and others inappropriate. Rather the knowledge in question is practical knowledge; it is knowledge of how one ought to act given the object that one is currently seeing.

As will come as no surprise by now, I intend to follow Merleau-Ponty. I will claim that it is by virtue of our possession of practical knowledge that our experiences have the phenomenology they do. I have said that our conscious experiences have their intentional features in virtue of their phenomenology. Thus it will follow that an experience has the intentional features it does in virtue of our practical knowledge.

I will argue for this thesis in two stages, corresponding to the two components in terms of which I have characterised an experience's phenomenology. Recall that I have said an experience's phenomenology is determined by (1) the objects and properties an experience seems to present and (2) the ways in which those objects and properties are presented in an experience. I will set about showing that each of these aspects of an experience's phenomenology can be constitutively determined by a subject's practical knowledge. I will look in part to the enactive theory of perception for an account of the practical knowledge in terms of which I shall account for an experience's phenomenology.

5. The Enactive Theory of Perception

Naturalising Phenomenology

According to the enactive theory, perception isn't something that just passively happens to us; it is something we, as perceiving creatures, do. Our perceptual experiences are something we act out through our movements. Perception is no longer understood as a process that takes place in the brain. The brain is of course a necessary condition for perceiving but it may not be a sufficient condition, as was argued above. Perceiving is instead taken to be a skilful activity of the whole animal which gets played out through the animal's interactions with its environment.

Alva Noë (2002 & 2004, ch.3) has suggested that we should think of the practical knowledge we draw on in perception as knowledge of how the appearance of an object will change with our movements. Each movement will bring about a change in what the object presents to us. As we move around a table, for instance, what we experience of the table will change with our movements. We can think of our visual perception of the table's shape as constituted by our knowledge of all the changes our experience of the table's shape can undergo as we move relative to it. Part of what it is for an experience to present what seems to be a table then is for us to know how the appearance of the table will change as we move around it. This practical knowledge is constitutive of our being presented with what seems to be a table in experience.

One objection one might raise against Noë's proposal is that it fails to do justice to Merleau-Ponty's point that the hidden sides of the object are present in experience albeit indeterminately. Doesn't Noë account for the experience I have of a thing's hidden sides in terms of the anticipations (the practical knowledge) I have about what I would see of the thing if I were to move around it? Didn't Merleau-Ponty show us that this description gets the phenomenology wrong misdescribing the hidden sides of a thing as sensibly absent when in fact they are positively present? The hidden sides of a thing

Naturalising Phenomenology

are positively present in my experience through the readiness or preparedness I exhibit to take hold of the thing in a certain way.

As an illustration of the phenomena Merleau-Ponty has in mind think about how my body is prepared and ready when I reach for the handle of a mug that is pointing away from me.³⁴ I relate to the mug through my taking hold of it, through the movements I direct towards it when I reach for it in order to take a drink. My doing so gives me a sense of the presence of its handle even though I am not currently seeing the mug's handle. Of course I can get it wrong and direct my reaching behaviour towards its handle when it has no handle. In this case I will have misrepresented the mug. The fact that I can misrepresent the mug in this way suggests that the handle was a part of my original experience, even though it wasn't any part of the view I had on the mug at the time.

What this objection shows us is that the knowledge which is constitutive of an object's seeming to be present in experience isn't just knowledge concerning the changes the appearances of a thing will undergo with our movement about it or its movement relative to us. We also have to account for the knowledge a creature has of the possibilities for action a thing affords. This knowledge of which actions are appropriate and which actions are inappropriate to the object of our experience can also contribute to what it is for an object to seem to be present in experience. This knowledge doesn't however look to be accounted for in terms of knowledge of sensorimotor contingencies.

Returning to the mug with its handle turned away from me, I am ready to deal with the mug before I direct my movement towards it. As Merleau-Ponty puts it, (discussing a different example) 'from the outset the grasping movement is magically at its completion.' (1962: 119) The hidden sides of a thing can be said to be present in my experience insofar as I am ready to deal

³⁴ This is a favourite example of Sean Kelly's.

Naturalising Phenomenology

with the thing, including its hidden sides, in my bodily interactions with it. I am, in other words, drawing upon what I have called my “operative understanding” of the thing when I experience the thing’s hidden sides. This operative understanding prepares me for my dealings with the thing, and enables me to undergo a perceptual experience of thing in its full bodily presence even though I am not presented with the thing in its entirety at any given moment.

The idea that what we perceive are the actions a thing affords is an idea we find in Gibson’s ecological theory of visual perception. Gibson (1979, ch.8) says that the environment of an animal affords opportunities for actions and it is these opportunities that we perceive on any give occasion. Gibson is insistent that visual perception works in such a way that we can pick up on a thing’s affordances directly and immediately. There is no need, Gibson claims, for an animal to reconstruct a scene in its brain based on information made available by the retinal image. Gibson claims instead that the structure of the light in an environment – what he calls “the ambient optic array” – uniquely specifies the layout of surfaces around the perceiving animal.³⁵

The environment the ambient optic array carries information about isn’t only composed of surfaces and objects. We also perceive the objects around us as affording us opportunities to do things. According to Gibson the ambient optic array exhibits certain systematic features. He distinguishes between “structural” and “transformational” invariants. The structural invariants in an ambient optic array are a consequence of certain stable and recurring features in the layout of an environment. Transformational

³⁵ The perceiving animal is stationed at a point in space at any give time and surrounding this point is light. The light surrounding the perceiving animal, Gibson calls ‘the ambient optic array’. The light is ‘ambient’ because it literally surrounds the perceiving animal. The ‘optic array’ is composed of the light that the animal’s eyes receive from surrounding objects. As I am sat here there are various objects arranged around me on my table, and straight ahead of me is a picture. All these objects are reflecting light some of which is received by my left eye, some of which is received by my right eye. This light is structured in such a way as provide me with information about the unchanging layout of my surrounding environment at this time. For further discussion of the ambient optic array see Gibson (1979, ch. 5).

Naturalising Phenomenology

invariants relate to ways in which the pattern of light changes with an animal's movements. Gibson claims that each structural invariant is an affordance. Whenever an animal perceives the layout of its environment, what it perceives are, in part, the opportunities an environment with this layout affords for its movement. When, for instance an animal perceives the brink of a cliff, what the animal sees is a surface one side of which can be walked on and the other side of which affords falling off and injury. It perceives such a surface because of a structural invariant in its optic array. Similarly a human being who perceives a red post box sees something that invites the posting of letters. The post box only has this significance, it only invites us to mail letters, because we belong to a community in which letters get written and are transported from one person to another by means of a mail service. Nevertheless, Gibson claims that as the kind of animal that inhabits an environment containing post boxes, we can directly perceive something as a post box. When we do so, what we have done is picked up on a structural invariant in the optic array of creatures like us.

If Gibson is right and we can directly and immediately pick up on the opportunities for action our environment affords this is only because we are drawing upon what I have called our operative understanding. It is our operative understanding for instance that enables us to know that red pillar boxes afford the posting of letters. Some of the simple examples Gibson offers like perceiving a tree to afford shelter, or to be climbable, or perceiving a surface to afford support, do not seem to involve much in the way of operative understanding. Nevertheless I want to say that to pick up on the information contained in the ambient array directly and immediately is to be drawing on one's operative understanding. It is only because one is in possession of this understanding that by perceiving invariant structures in the ambient optic array one can also perceive opportunities for action.

Naturalising Phenomenology

Noë (2004: ch.3, §9) has offered the following helpful suggestion about how to interpret Gibson's claim that the ambient optic array carries information that specifies the layout of one's environment:

'The ambient optic array, simply put, is how things look from here in these conditions. The sense of the Gibsonian claim that the ambient optic array specifies the environment (unlike the pattern of irradiation on the retina), is that how things look from here in these conditions specifies how they are, or rather, it does so for a suitably knowledgeable animal, one in possession of and ready to apply sensorimotor skill.' (2004: 104)

This seems to me exactly right. The claim is threefold.

(1) What Gibson calls the structured invariants present in the ambient optic array can be identified with how the environment appears to us at any given moment.

(2) How the environment appears to an animal at any moment specifies how things are in its local environment. This is a statement of the phenomenal intentionality thesis according to which an experience's phenomenology is constitutive of its intentional features.

(3) An animal's possession of practical knowledge (what Noë refers to as 'sensorimotor skill') determines how the environment appears to an animal on any given occasion.

Noë errs slightly in his characterisation of this practical knowledge. His description of the practical knowledge we draw upon in perception doesn't really bring out the ways in which our embodiment can shape our perception of an object and its properties.

Noë characterises our perception of affordances in terms of the possibilities for movement that a thing affords. 'To perceive is (among other things) to learn how the environment structures one's possibilities for movement...' (Noë, 2004: 105) He goes on to mention in passing that we also perceive a thing as affording possibilities for action. However the

Naturalising Phenomenology

account he gives of practical knowledge in terms of knowledge of sensorimotor contingencies only tells us what it is to perceive the opportunities for *movement* a thing affords. I shall argue that movement is only one of a class of actions we can perceive a thing to afford.

Noë characterises knowledge of sensorimotor contingencies as knowledge of the systematic ways in which how things look, sound, smell or feel depends on one's movements. There are sensorimotor contingencies that specifically relate to our sensory apparatus. These include the fact that 'when the eyes rotate, the sensory stimulation on the retina shifts and distorts in a very particular way, determined by the size of the eye movement, the spherical shape of the retina, and the nature of ocular optics.' (O'Regan and Noë, 2001: 941) We know what effects our eye movements have on what we are seeing. This constitutes one kind of sensorimotor contingency of which we have knowledge. We also know that 'the flow pattern on the retina is an expanding flow when the body moves forward, and contracting when the body moves backward.' (op cit) This is another example of a kind of knowledge we have about how our movements can effect what we see.

Another set of sensorimotor contingencies relate to the visual attributes of the objects we are sensing. The visual quality associated with a thing's shape is, O'Regan and Noë suggest 'precisely the set of all potential distortions that the shape undergoes when it is moved relative to us, or when we move relative to it. Although this is an infinite set, the brain can abstract from this set a series of laws, and it is this set of laws which codes shape.' (O'Regan and Noë, 2001: 942) They go on to add 'the structure of the laws abstracted from the sensorimotor contingencies associated with flat, concave, and convex surfaces, corners and so on, will be a neural-code-independent indication of their different natures.' (op cit)

What we perceive on any given occasion is on this view the result of our drawing on knowledge of how our sensory experience of a thing will change

Naturalising Phenomenology

with our *movements*. This only partially characterises the practical knowledge we are drawing on in virtue of which things seem to us the way they do. In addition we have a knowledge of which actions are appropriate to the object of our perception and which are inappropriate. When we see a heavy box which we are about to lift for instance we prepare our body to deal with the box differently from how we would be ready to deal with the box if we perceived it to be light.

As another example of the phenomena I have in mind consider Milner and Goodale's patient D.F.³⁶ Despite D.F.'s not being able to visually identify an object's features she is nevertheless 'capable of responding differentially' to features like an object's size, shape and orientation. She is able for instance to post a card, or insert her hand through a slot set at different angles, even though she cannot report on the angle of the slot. The understanding D.F. has of the slot is a bodily understanding. She understands the slot only through her dealings with it.

I claim that the understanding of a thing we have through our embodiment also plays a constitutive role in determining the ways in which a thing appears to us. We don't only perceive the possibilities for movement a thing affords. We also perceive its possibilities for action in a much wider sense. Crucially we can perceive the opportunities for action a thing affords only because we are in possession of a certain kind of practical knowledge. What we learn from Merleau-Ponty (with Sean Kelly's help) is that there are bodily ways of understanding an object which contribute towards this practical knowledge. These bodily forms of understanding also contribute to determining the content of our perceptual experiences.

I will conclude this section by considering how the second aspect of an experience's phenomenology, the modes of presentation characteristic of our perceptual experiences, might be determined by our practical knowledge.

³⁶ I am indebted here to Kelly's discussion of D.F., see his (2002) & (2003).

Naturalising Phenomenology

Consider as an example of this second aspect of an experience's phenomenology, perceptual constancy phenomena. Suppose I am looking at a wall painted mint green. What my experience represents is a wall coloured green, yet some parts of the wall appear darker than other parts. I say that these disparities in the appearance of the wall are due to the wall being presented differently under different lighting conditions. Some parts of the wall are better lit than others and this is why some parts of the wall appear lighter than other parts. Now the puzzle that this kind of phenomena raises is how it can be that we see a wall that is uniformly coloured mint green when some parts of the wall appear to be darker than other parts. How is it that the wall can be perceived to have a constant colour when some parts of it appear different from other parts?

Noë (2004) characterises this problem in the following terms: 'We experience the presence of a uniform colour which, strictly speaking, we do not see. Or rather, the actual uniform colour of the wall's surface is present in perception amodally, it is present but absent, in the same way as the tomato's backside, or the blocked parts of the cat.' (2004: 128)

Again Noë takes the solution to this puzzle to lie in our knowledge of sensorimotor contingencies. Noë distinguishes between two kinds of sensorimotor contingencies both of which he takes to be relevant to determining our experience of a thing's apparent colour. The first kind relate to 'The way sensory stimulation is affected by changes in a perceiver's geometrical relation to an object...the way stimulation varies as a result of the perceiver's manipulation of an object (e.g. turning it in relation to a light source).' (Noë, 2004: 129) The second kind relate to 'patterns of dependence between sensory stimulation and the *object's* movement, or the object's changing relation to its surrounding.' (op cit, p130)

Noë's answer to the puzzle of how we can perceive a wall to have a uniform colour even though this is not what we are presented with in

Naturalising Phenomenology

perception, appeals to the knowledge we have of sensorimotor contingencies of these two types. A thing's apparent colour can vary in all kinds of ways some of which have to do with our movement, others of which relate to variations in lighting conditions and in the colour of the background and surrounding objects.³⁷ Noë hypothesises that nevertheless we have knowledge of how all these factors combine to determine an object's apparent colour at any moment. Our understanding how a thing's apparent colour will change as viewing conditions change explains how we come to have experiences which present us with coloured things, Noë suggests.

Again it seems to me this is only part (albeit an important part) of the story. Once again we can look to Merleau-Ponty via Kelly (1999& 2004) to supply the rest of the story. Merleau-Ponty thought that the lighting conditions form a part of the context or the background against which a thing appears coloured. He insists that the lighting conditions can play a positive role in determining the kind of experience we enjoy. However it contributes to the way we see things only indeterminately.³⁸ For Merleau-Ponty (according to Kelly) the lighting context enters into our experience through our knowledge of how *best* to see an object. We know where to move our eyes in order to achieve the best view of an object's colour. We know for instance that dark colours are seen best in bright light, and that bright colours are seen best in dimmer light.³⁹ What is it for us to see the wall as coloured mint green all over? Merleau-Ponty will say it is for me to know given the current lighting conditions, where I would need to look in order to gain the best view of the object's colour.

Again Merleau-Ponty appeals to a kind of knowledge that is essentially related to my embodiment to explain colour experience. In this case I have

³⁷ See Noë (2004: 132)

³⁸ Notice the similarity with the position Merleau-Ponty has taken on a thing's hidden sides – these also enter into experience only indeterminately, by means of the experience we have of them by virtue of our embodiment.

³⁹ Kelly (2004: 86)

Naturalising Phenomenology

knowledge of how the light would have to change in order to see the colour of the thing best of all. Kelly (2004) explains the idea as follows:

‘To speak mathematically, I experience the light not as a determinate quantity but in terms of the direction, and perhaps even the slope, of the improvement curve. If we think of the improvement curve as the curve that measures the quantity of light against the quality of the viewing conditions, then what my experience tells me at any given moment is whether more or less light will improve my view, and also perhaps, how drastic the improvement will be. In this way the lighting plays a positive role in my experience but is never registered determinately.’ (Kelly, 2004: 85-6)

It is an interesting question which I cannot enter into here in any detail, the extent to which Noë and Kelly have offered competing explanations of the kinds of practical knowledge which might determine the contents of my colour experience. Noë’s position looks to me to be not dissimilar to what Husserl might have said on this question. Recall how Husserl took the hidden sides of a thing to be present in experience only in the form of hypotheses about what we would see if we were to change our relation to a thing. Noë’s position on our perception of a thing’s colour looks to me to involve a similar claim. Noë seems to be suggesting that an important part of what it is to experience a thing as having a uniform colour involves us drawing on our knowledge of sensorimotor contingencies so as to form hypotheses or expectations about how a thing’s colour will change under different viewing conditions.

If this is the correct way of reading Noë, it would seem that we do indeed have two competing accounts. Husserl and Noë say that what we might call the perceptual context enters into perception only by contributing to the hypotheses we frame about how our perceptual experiences of a thing will change with our movements. Merleau-Ponty and Kelly say that the lighting context can enter into perception through our bodily understanding of under what lighting conditions the object can be best seen.

Naturalising Phenomenology

I cannot attempt to decide between these two explanations here. It suffices to note that there are two hypotheses here about how our experiences get their phenomenal intentionality. Both hypotheses can be understood as empirical hypotheses. Thus I take myself to have presented the beginning of a story about how the naturalist might go about assimilating the phenomenologist's idea of phenomenal intentionality. I have proposed an account of phenomenal intentionality which appeals to certain kinds of practical knowledge to account for why the environment appears at it does to an animal on any given occasion. The idea is that being in possession of certain kinds of practical knowledge is what makes it the case that the world appears to us at it does. Admittedly there is more work to be done in explaining just how it is that being in possession of this practical knowledge suffices to make it the case that our experiences present the world as appearing thus and so. This is a task for my future research. On the supposition that possession of this practical knowledge is something that can be given a functional-physical description I take it I have sketched the beginnings of a story about how a naturalist might go about integrating phenomenal intentionality into her thinking about minds.

In the next section I will explain how possession of practical knowledge does look like something that can be explained in functional-physical terms. I will begin by connecting the account I have presented of phenomenal intentionality thus far with the conclusion that there are subjective facts. We will see that the existence of phenomenal intentionality requires us to revise our thinking about the nature and function of perception. However this revised conception of the function of perception finds support in recent work on perception in cognitive science and neuroscience. Thus the shift in our thinking about perception required if we are to admit subjective facts, is a shift already taking place within the science of mind.

Naturalising Phenomenology

6. Why Perception isn't Detection

Once we concede that there exists a kind of intentionality which is constitutively determined by a subject's practical knowledge in the way I have just described, it is a short step to the conclusion that there are subjective facts. The practical knowledge I have invoked above is nothing other than what I have been calling "operative understanding". In chapter 2 we saw how it is in virtue of our operative understanding that we are able to perceive a thing as calling for or drawing from us certain actions. To perceive an entity in this way is precisely to understand that entity as an entity of a certain kind. It is to understand the entity as an entity that calls for certain actions.

Consider Heidegger's hammer as an illustration of this point. In order to perceive a thing as a hammer, as a thing which can be used to fasten together two pieces of wood say, one must draw on one's understanding of how a hammer is to be used. One must represent the hammer in a certain way, in a way that is determined by one's operative understanding of hammers. Now consider someone who hasn't the faintest idea about how to use a hammer or what hammers are for. This creature won't perceive the hammer in the way that we perceive it: the creature won't perceive the hammer in terms of how it is to be used. By hypothesis she doesn't know how hammers are used. Thus she cannot represent a thing as a hammer. It follows then that there are things that can be represented only by a creature with the requisite operative understanding.

Someone might object that this doesn't establish the existence of things that can be represented only given the right operative understanding. It only establishes the existence of ways of representing that require possession of an operative understanding. The creature can perfectly well perceive something that is a hammer, assuming her senses are all functioning normally. What she cannot do is represent the thing she perceives as a

Naturalising Phenomenology

hammer. What she lacks is the recognitional concept for hammers, and this is why she cannot perceive the thing as a hammer.

By way of an initial response to this objection, I would reply that she not only lacks the ability to recognise hammers. She doesn't know what to do with hammers. This is why the existential phenomenologist claims that she doesn't perceive something that is a hammer, something that has the being of a hammer. A longer response requires us to go back to the notion of operative understanding introduced in chapter 2. I introduced the notion of an operative understanding in the course of describing the way of existing characteristic of persons which Heidegger calls "being-in-the-world". Operative understanding was taken to be one of the structures constitutive of our way of existing as persons. This is to say that operative understanding forms a part of what it takes for an animal to be a being-in-the-world. When a creature has the required operative understanding she comes to inhabit a world, a world that contains hammers, for instance, and the other tools of trade characteristic of carpentry. Before she acquires the necessary operative understanding she doesn't live in a world in which anything can appear to be a hammer. She doesn't live in a world in which there are hammers because things of this kind have no significance for her.

For the remainder of this section I want to consider how the latter idea finds support from within naturalism. The claim to repeat is that our ways of existing as persons can literally shape the world we inhabit. Let us begin by contrasting this idea with a rival view of the relation between an animal and its environment. This rival view finds its clearest articulation in a natural way of thinking about perception. According to this view the senses function to provide the brain with information about what is going on in the animal's environment (or in the animal's body, in the case of proprioception). On this view a brain state is about a thing x in virtue of some relation between the animal whose brain it is, and some event in the animal's environment. The

Naturalising Phenomenology

senses are taken to be the brain's window onto the world. The function of the senses is simply to provide the brain with information about what is where.

This view is based on the reasonable claim that were the senses not playing this role for the brain, the animal wouldn't know how things stood either with respect to its local environment or in its own body. I shall call this the "detection view" of perception, so-called because it construes the function of the senses to faithfully report back to the brain how things are in the world external to the brain.

The detection view of perception conflicts with the story I have been telling according to which what we perceive are opportunities for action. On this view the job of the senses isn't simply to report back to the brain about what is where. In addition the senses tell an animal what it should do. This is a job the senses can perform because it is a part of the animal's way of existing to know how to find its way about in the world. This is to say that the animal knows how to make use of the things it perceives to achieve its goals.

Kathleen Akins (1996) has presented a similar view of the senses in attacking what I have called the detection view. Akins characterises the senses as fundamentally "narcissistic"; the animal is constantly posing the question 'how does all of this relate to me?' and it is the job of the senses to deliver an answer to this question. Akins gives a nice example to illustrate her proposal, she asks us to think about the function of thermoreception. Given the detection view we would expect the thermoreceptive system to be answering the question 'what is my skin temperature at x' where 'x' picks out a particular region of my body. However it turns out this fails to capture what our thermoreceptive systems do.

There are four different types of thermoreceptors, two of which respond only to extreme conditions of very high and very low temperatures causing in us sensations of pain. When the conditions are not in this way extreme, we rely on thermoreceptors for our sensations of warmth and cold. Akins begins

Naturalising Phenomenology

by noting that we have many more cold receptors than warm receptors. Moreover, some parts of the body have more receptors than other parts, and the ratio of cold to warm receptors likewise varies from one part of the body to another. Akins goes on to note that how each receptor responds to a temperature change will depend upon the starting temperature. When a part of the body is already cold, applying a cold stimulus to this part of the body – say you run your hands under a cold tap by accident – will evoke a dramatic response. A warm stimulus will however evoke a gradual decrease in activity in the cold receptors. The reverse is true of warm receptors. If you put already warm hands under a hot tap the warm receptor will burst into life, a cold stimulus will produce a gradual decrease in activity.

What bearing does this have on the detection view of the senses? First it doesn't look right to say that the function of thermoreceptors is to report back to the brain, the temperature of particular bodily regions. One and the same bodily temperature can give rise to a variety of sensations. Akins makes the point as follows: 'thermal sensations are a function of the firing rates of a neural population and because the absolute number and ratio of the two different receptors (the cold and warm receptors) differ from one part of the body to another, exactly the same skin temperature can give rise to a variety of sensations.' (Akins, 1996: 351, my addition in brackets) Furthermore, the response that does occur in a particular part of the body at a particular time will depend on the starting temperature in that part of the body. 'The felt change in temperature for a specific temperature change will depend upon the starting temperature of the skin. If the temperature of a warm spot is increased at the bottom of its response range, the dynamic burst (very roughly, the firing rate of neurons) will be very small; if it is warmed at the top of its response range, the burst will be very large.' (op cit, my addition in brackets)

Naturalising Phenomenology

On the basis of points like these, Akins mounts a convincing case against a detection view of thermoreception. She establishes that the function of thermoreception cannot be to record absolute temperature or temperature change at bodily regions. Akins sums matters up nicely:

'What the organism is worried about, in the best of narcissistic traditions, is its own comfort. The system is not asking, 'What is it like out there'? – a question about the objective temperature states of the body's skin. Rather, it is doing something – informing the brain about the presence of any relevant thermal events. Relevant of course to itself.' (Akins, 1996: 349)

When we are considering the function of perception the question to ask is not what the senses ought to be detecting. The question we should ask is rather what should the senses be doing, where what the senses should be doing will be decided by the interests of the animal at that time. What the animal is doing will tell us what it needs to be representing. Of course to know what implications the things it is sensing have for action requires skill on the part of the animal. It will require the animal to draw on its operative understanding. Thus at first glance Akins' narcissistic view of the senses would seem to fit comfortably with the view of perception I have been presenting from existential phenomenology.

The existential phenomenologist however advances a further claim which doesn't obviously find support in Akins. They claim that not only is it the case that an animal perceives what it is in its interest to perceive, where those interests will be dictated by what the animal is engaged in doing at the time.⁴⁰ They claim in addition that the objects an animal perceives have an existence

⁴⁰ Akins' point seems to be that we must understand the function of perception in the context of the behaviours it services. She notes a little later in her paper (p.354) the significant connections there are between the parts of the mammalian brain dedicated to seeing and other parts of the brain dedicated to motor behaviour. This provides important support for the idea of perception as being for the selection of action.

Naturalising Phenomenology

which is *shaped* by the animal's interests. I will argue next that this is an idea that also finds support in psychology and neuroscience.

Akins says: 'Each and every sensory system, no matter how sophisticated or simple, is tied to a set (sometimes a very large set) of behavioural tasks. No matter what else the senses do, in the end, they must inform movement or action.' (1996: 352) There are at least two ways of understanding the claim that each sensory system is tied to set of behavioural tasks. The first reading takes the tie to be instrumental. The information supplied by each sensory system is, on this understanding, made available to other systems that control and generate motor behaviour. The close ties between perception and action on this understanding, just consist in perceptual information being used directly to control motor behaviour.

There is however another more radical way of understanding the close ties between perception and action. On this understanding, perception just is a kind of activity, a probing or exploring of the environment. On this understanding the connection between perception and action isn't just instrumental; the connection is constitutive. What it is for an animal to sense its environment is just for the animal to engage in a certain kind of activity, the activity of exploring the environment by means of its senses. This is the view we find defended in existential phenomenology. The existential phenomenologist takes there to be a constitutive connection between perceiving and acting. He takes perception to be a mode of skilful engagement with the world.

Now suppose we grant that perceiving is an activity in which the animal engages. It follows that the objects of perception will always be *understood* by the animal in terms of the implications they have for movement and action. This is to say that the ways in which an object figures in our perceptual experience will be partially determined by the animal's operative understanding. The claim that the connection between perception and action

Naturalising Phenomenology

is constitutive would seem to imply that the objects of an animal's perceptual experience are indeed shaped by an animal's operative understanding.⁴¹

This claim that perceiving is a kind of action is one we find outside of existential phenomenology. It is supported by work in cognitive science, neuroscience, and robotics research as well, and has recently been defended by a handful of philosophers of perception.⁴² It is a central idea in both Gibson's ecological theory of perception and in the enactive account of perception. Ballard (1991) has suggested that 'vision is best understood in the context of the visual behaviours that the system is engaged in' (1991: 57). So far this is compatible with an instrumental understanding of the relation between vision and action, where vision is made use of, perhaps directly, in the control of behaviour. However this is not what Ballard has in mind. He makes clear that the behaviours he has in mind are integral to the activity of seeing itself. Vision as he describes it is active and exploratory. In particular he emphasises the role that gaze control has in visual processing, where gaze control is a collection of processes that enable us to keep the fovea of the eye fixed on a particular target while either we move or the thing our gaze is fixed upon moves. Ballard's account of vision doesn't eschew the role of information processing in vision as did Gibson and his followers. Rather

⁴¹ Someone might be willing to grant that there is a constitutive connection between perception and action but not be willing to grant that our perceptual experiences shape the objects of our experience. He might say that the object as it is represented is shaped by our ways of existing so there is a constitutive tie between our ways of representing the world and our ways of being. He might however add that there is no such tie between the object itself and our existence.

This is a difficult point which we have found ourselves returning to repeatedly. I don't have anything new to add which I haven't already said elsewhere in response to this kind of objection. I would say however that the weaker claim the objector concedes is enough to establish the existence of subjective facts, as has already been noted earlier (see the final section of chapter 3). For to concede that there are objects or facts or whatever that can only be represented by a creature that has the required operative understanding, is to concede that there are facts which can only be represented from a subject's point of view, that is to say, from the point of a subject that has the right kind of operative understanding.

⁴² For a review of work in cognitive science, robotics and neuroscience which supports this picture see Clark (1997 and 2000: ch.'s 5 & 6). Also see Hurley (1998: ch.'s 9 & 10); Haugeland (1998, ch.9); Noë (2004, ch.'s 1 and 2); Rowlands (1999, ch.5 and 2003, ch.10); Thompson (1995) and Thompson and Varela (2001).

Naturalising Phenomenology

Ballard gives an account of vision where use is made of the animal's behaviour to simplify the kinds of computations its visual system performs.

Rodney Brooks' work in robotics also suggests a view of perception as constitutively tied to action. Brooks' robots are composed of subsystems or what he calls "layers" each of which is dedicated to carrying out a particular activity or task such as detecting and avoiding obstacles, wandering around, scanning surfaces etc. Each layer operates independently of the other layers, the overall behaviour of the robot being under the control of one or the other of its layers. To understand what takes place in each layer we have to consider what the layer is doing. What the layer is dedicated to doing will often require the use of perception. However the contribution made by perception cannot be understood apart from the activity each layer is dedicated to carrying out. Each activity involves a pattern of interactions with the world. Brooks says in discussing a simple creature built to avoid obstacles:

'...there need be no clear distinction between a "perception subsystem", a "central system" and an "action system". In fact, there may well be two independent channels connecting sensing to action (one for initiating motion, and one for emergency halts), so there is no single place where "perception" delivers a representation of the world in the traditional sense.'
(Brooks, 1991: 144)

Consider the layer in a robot of this kind dedicated to initiating motion. Clearly it has to be making use of what it perceives if it is to avoid obstacles. However the use the creature makes of perception cannot be understood apart from what the robot is doing, which is initiating motion, in the case of the layer we are considering. Thus Brooks is a clear example of someone whose

Naturalising Phenomenology

empirical work is predicated on the assumption that perception and action are constitutively tied.⁴³

As a final example, consider the account Thompson (1995 & 2000) gives of the function of colour perception. Thompson wants to steer a middle course between views of colour perception which take its function to be the detection of surface reflectance properties, and views which take colour to be something the brain projects onto the world. Thompson points out that there are significant differences in colour vision throughout the animal kingdom. Normal human perceivers are trichromats, which is to say that they possess three different types of cone photoreceptor. Our colour space – the space of different colours we can discriminate between – is three-dimensional. Birds, fishes, amphibians and reptiles are by contrast tetrachromats, they possess a fourth type of photoreceptor. These animals enjoy experiences of colour which are tetravariant, having another hue dimension in addition to ours. Thompson hypothesises that given this difference in the visual apparatus of these creatures, they will be able to perceive novel hues we cannot see.

Thompson takes these differences to indicate that there is no common property which vision in these different animals has the function of detecting. Thompson agrees with Akins that when we are thinking about the function of colour vision the question to ask is “what does colour vision do for the animal in its environment.” (Thompson, 2000: 165)⁴⁴ Thompson comes up with the following answer:

‘...the primary role of colour vision may be to enact a perceptual quality space that integrates physically heterogeneous environmental properties into a small number of perceptual equivalence classes. These relatively

⁴³ In fact Brooks willingly endorses the idea that an animal can shape the environment it lives in. Commenting on work in AI which attempts to model special purpose problem solving he has this to say: ‘...as Uexküll and others have pointed out, each animal species, and clearly each robot species with their own distinctively non-human sensor suites, will have their own *Merkwelt* (perceptual world).’ (Brooks, 1988: 141)

⁴⁴ He cites a co-authored paper Thompson, Palacios & Varela (1992) in support of this claim.

Naturalising Phenomenology

stable perceptual categories would facilitate the identification of aspects of the environment and then guide behaviour accordingly.’ (Thompson, 2000: 182)

Thus for Thompson the purpose of colour vision isn’t to detect physically invariant surface reflectance properties. It is to make available to the perceiving animal, stable perceptual categories that can be used to meet the animal’s needs in the course of its perceptually guided activities. So it turns out that the same thing can have different colours – it can have one colour for me and a different colour for a tetrachromatic pigeon, say.⁴⁵

I conclude then that the argument from the contents of experience for subjective facts is one that a naturalist could buy. A commitment to subjective facts follows from the empirical hypothesis that action and perception are constitutively connected. Whatever empirical work supports this hypothesis, I claim will also support the existence of subjective facts. For perception and action can be constitutively tied only if the animal is drawing on practical knowledge which tells it what the implications for action are of the object it is perceiving. A creature that makes use of its practical knowledge in this way is, I claim, a creature that lives in a world shaped by its existence.

The remainder of my chapter will take up the second reason that was given for introducing subjective facts. This was the phenomenologist’s claim that every conscious experience is also pre-reflectively self-conscious. It will turn out that we are already well on our way to having an answer to this question. For I shall argue eventually that the kind of operative understanding that has been appealed to in developing an account of phenomenal intentionality can only be had by a creature that has access to its conscious experiences. A creature that has this kind of access is pre-reflectively self-conscious.

⁴⁵ See Thompson (2000: 183) for more on this point.

Naturalising Phenomenology

7. The Feeling of Ownership

I will begin with a brief reminder of the thesis to be defended – the claim that every conscious experience is also PRSC. This thesis was introduced in answering the question of what it is for an experience to be phenomenally conscious. An experience is phenomenally conscious when there is something the having of the experience is like for a subject.

I claimed that there is something an experience is like for a subject when the occurrence of this experience is itself experienced by the subject. There is something the experience is like for its subject when that very experience is something the subject experiences. Consider my experience of the blue sky: this experience presents the sky to me as seeming a certain way only if the experience is itself experienced by me. When I do not experience my experiences they do not seem any particular way to me: they lack a subjective character.

Whenever an experience is itself experienced by a subject, I shall say that this experience is accompanied by a *feeling of ownership*. The feeling of ownership accompanies an experience when a subject has a sense of himself having that very experience. A subject's experiencing himself having an experience is an experience of ownership. For "to have" and "to own" are different ways of describing one and the same relation. Thus, when a subject experiences himself having an experience, he also experiences himself as the owner of the experience that is happening to him.

Now I have said that there is something an experience is like for a subject only when the experience is itself experienced. Thus, we can say that there is something an experience is like for a subject only if it is accompanied by a feeling of ownership.

An experience which is accompanied by a feeling of ownership is an experience that is had self-consciously. The kind of self-consciousness that is characteristic of the feeling of ownership should be distinguished from what

Naturalising Phenomenology

we might call “reflective consciousness”. A subject need not introspect or reflect on himself in order to be consciously aware that he himself is having an experience. What I am calling the feeling of ownership accompanies our experiences without us actively or deliberately doing anything. Merleau-Ponty makes the point well:

“At the root of all our experience and all our reflections, we find...a being which immediately recognises itself, because it is its knowledge both of itself and of all things, and which knows its own existence, not by observation and as a given fact, nor by inference from any idea of itself, but through direct contact with that existence.” (Merleau-Ponty 1962:371)

Thus we need to distinguish between a subject’s being *reflectively* conscious of himself having an experience and what I labelled in chapter 5, a subject’s *pre-reflective consciousness*. It is pre-reflective and not reflective consciousness that must be explained if we are to explain what it is for an experience to have subjective character.

Pre-reflective consciousness is a kind of *peripheral* or *implicit* consciousness a creature has of itself. Consider the sentence ‘Smith is conscious of her thinking that she herself is feeling nervous’.⁴⁶ This sentence reports a state of affairs in which Smith is conscious of herself having a particular thought. Moreover, the thought she is having about herself is a distinct state of mind from the feeling of nervousness she is experiencing. The thought the subject has about herself takes as its object the feeling of nervousness. Now contrast this state of affairs with that reported by the following sentence: ‘Smith is self-consciously feeling nervous’. Here, the state of mind reported is not distinct from Smith’s feeling of nervousness. Rather, Smith’s self-consciousness *modifies* the feeling of nervousness she experiences.

⁴⁶ Here I am drawing on a useful distinction Kriegel (2003 and 2004) makes between *transitive* and *intransitive* self-consciousness.

Naturalising Phenomenology

Going back to Brentano, we can say that Smith is *primarily* conscious of her feeling of nervousness and she is only *peripherally* or *secondarily* aware of herself having this experience. Thus, the difference between reflective and pre-reflective consciousness concerns the way in which the subject is self-conscious. When a subject is reflectively conscious of herself having an experience, she is *attentively* or *focally* aware of herself and the state of mind she is in. Whereas, when a subject is pre-reflectively conscious she is only *inattentively* or *peripherally* aware of herself having an experience.⁴⁷

A subject is pre-reflectively self-conscious when her experiences are accompanied by a feeling of ownership. If this is right, a naturalistic answer to what I have earlier called “the experience of content question” – the question why there is something rather than nothing it is like to undergo a conscious experience – will lie in an explanation of what it is for an experience to be accompanied by a feeling of ownership.

The feeling of ownership is such that the subject does not experience himself simply as a material thing existing alongside other material things, but instead experiences himself as a material thing that is the owner of various experiences. The subject experiences himself *qua* subject. If we are to explain what it is for an experience to be accompanied by a feeling of ownership, some explanation must be given of what it is for a creature to represent itself *qua* subject.

8. Towards a Naturalistic Account of Pre-Reflective Self-Consciousness

In chapter 6 I introduced a variety of representational state which I claimed could make a subject implicitly aware of herself *qua* subject at the same time as making her explicitly conscious of something in her environment. What I shall set about doing in what remains of this chapter is explaining just what it is for a subject to be implicitly aware of herself *qua* subject. I will suggest that

⁴⁷ I am indebted to Kriegel (2004) for this way of formulating the distinction.

Naturalising Phenomenology

we can understand this awareness as part and parcel of a subject's exercising her operative understanding. PRSC and operative understanding form parts of a single package. Indeed I shall argue that a subject couldn't possibly perceive an object as having implications for her actions if she weren't in this way conscious of herself *qua* subject of this experience.

I characterised PRSC as a form of bodily self-awareness in chapter 6. I said that a subject is continuously aware of her body *qua* subject of her experiences. Let us begin then by clarifying just what this continuous awareness consists in. It would be a mistake to understand it as a form of perceptual awareness. Perceptual awareness is intentional – it is awareness directed towards an object. Yet PRSC isn't intentional. A subject's body isn't among the objects he perceives whenever he undergoes a perceptual experience, though of course it might be on some occasions. I claim instead that a subject is aware of her body as that *with which* she perceives things. To give an example, consider the kind of awareness that accompanies one's experience of running one's fingers across a piece of cloth. When one moves one's fingers across the cloth, one's fingers do not figure among the object of which one is aware, unless one deliberately attends to them. You will nevertheless be aware of your fingers as that with which you are making contact with the cloth.

I suggest then that we think of PRSC as an awareness of one's own body as that with which one perceives. This characterisation of PRSC can help us understand why it should be that the feeling of ownership accompanies each of our experiences. The idea is that whenever a conscious perceptual experience occurs we are aware of our body as that with which we are perceiving things. Why should this kind of awareness form a constant backdrop to all of our experiences? I want to finish up by suggesting that it does so because of the role that our practical knowledge plays in constituting the appearance of things.

Naturalising Phenomenology

When one is behaving skilfully one has a sense of which actions are appropriate to what one is doing and which are not. One selects from amongst a number of possible actions which of them is appropriate to the task at hand. It is characteristic of skilful behaviour that one doesn't need to give any thought to what one is doing. One simply knows how to act. One is ready to deal with things, and this readiness consists in an important part in what I have earlier characterised as a bodily understanding we have of the things we perceive. For the most part then my body is ready to take hold of the things with which I deal. This kind of readiness my body exhibits at any given moment is, I have argued, an important determining factor in why things appear to me as they do at any given moment.

Now I want to suggest that my body can exhibit this kind of readiness to deal with a situation appropriately only because I am aware of my body as that with which I perceive things. Consider what happens when I play a tennis serve. To what extent are the movements I make something for which I can take credit? They are certainly not something I know *that* I am doing as I make them. The serve happens too quickly (or it would if I was any good at tennis) for me to know that I am making certain movements. Still these movements are under my control. They are not under my reflective control: they are unreflective or as I have been putting it "pre-reflective". The tennis serve is under my control in that there is something my movements are directed towards bringing about, namely getting the ball to land in the right part of the court. Insofar as this activity is under my control, and I take it that it will be under my control if I am a skilled tennis player, I must be conscious of it. The consciousness I have of it is precisely a consciousness of my body as that with which I am acting. It is through this consciousness I have of my body that I am ready to move in ways appropriate to taking a tennis serve.

I conclude then that the two questions I have been attempting to answer about the nature of consciousness meet up in the notion of practical

Naturalising Phenomenology

knowledge. My experiences have the intentional content they do in virtue of my practical knowledge. I am able to exercise this practical knowledge only because I am PRSC. PRSC gives me a kind of access to what I am experiencing in virtue of which I am ready to deal with the familiar things which fill my environment. Now clearly there is much more to be said here, but I have at least begun to sketch an answer to the question of what it is to be a conscious creature.

Naturalising Phenomenology

Conclusion

I have been arguing that naturalistic accounts of the mind have run into an explanatory gap because they have failed to recognise the existence of subjective facts. If naturalists are to close the explanatory gap it will only be by making room for subjective facts. My aim in this chapter has been to take up the two arguments that have been given in support of subjective facts, and show how they can be made to fit within a naturalistic account of the mind. In doing so my aim was to show exactly how a naturalist could go about admitting the existence of subjective facts. I take it that I have given an answer which points in the right general direction one might head in, were one to set about naturalising phenomenology. More significantly still I have shown that the phenomenologist's opposition to naturalism is unfounded. For I have shown that the naturalist can easily take on board a good many of the key claims the phenomenologist makes in describing the structure of our conscious experiences.

There is still much work to be done of course. Central to the account I have sketched is the idea of our practical knowledge or operative understanding accounting for phenomenal intentionality. If this account is to get off the ground, I must explain in much more detail just how it is that being in possession of the right kind of practical knowledge can have as a consequence that the world comes to appear in certain ways. There is also much more to be said about the account I have sketched, all too briefly, of PRSC in the final sections of this chapter. One question I was unable to pursue but would have liked to have examined is the relation between this PRSC and certain primitive forms of self-awareness which psychologists in the Ecological tradition such as Neisser (1988) have introduced, and which philosophers have recently begun discussing under the banner of nonconceptual self-consciousness (see for instance Bermudez (1998)).

Naturalising Phenomenology

Another intriguing avenue I would like to explore in the future is the relation between PRSC and our consciousness of time. Husserl claimed that every conscious experience has a temporal structure such that it contains a retentive reference to past moments of experience, an openness to the present, and a protentive anticipation of moments of experience that are about to happen. The idea that consciousness has this temporal structure was to prove immensely important for the existential phenomenologists too. Heidegger for instance was to argue that temporality forms the very structure of our being-in-the-world. Thus there is a significant theme here in phenomenology which connects with many of the ideas I have introduced but which I have been unable to tackle within the confines of this project.

Even more intriguing is the possibility of developing a naturalistic account of the claim that our experience, and maybe even our very existence, has this temporal structure. Varela (1999) for instance makes use of data from neuroscience to develop a model of how an experience could have this retentive-protentive structure. While Van Gelder (1999) shows how this idea fits with a dynamicist approach to thinking about mind in recent cognitive science.

There are many large questions which remain to be pursued. Still I have begun the work of developing an account of mind which allows us to understand appearances as essentially subjective as well as taking them to be a part of the fabric of reality. Now that the phenomenologist's arguments against naturalism have been revealed to be without substance, the way has been cleared to begin to develop a naturalised account of phenomenology.