THE CONTENT OF PHENOMENAL STATES: IN DEFENSE OF A MIXED THEORY

THE CONTENT OF PHENOMENAL STATES

IN DEFENSE OF A MIXED THEORY

By

JAYAR LA FONTAINE, B.A.

A Thesis

Submitted to the School of Graduate Studies

in Partial Fulfilment of the Requirements

for the Degree

Master of Arts

McMaster University

MASTER OF ARTS (2006) McMaster University
(Philosophy) Hamilton, Ontario

TITLE: The Content of Phenomenal States: In Defense of a Mixed Theory

AUTHOR: Jayar La Fontaine, B.A. (Brock University)

SUPERVISOR: Professor B. Garrett

NUMBER OF PAGES: vii, 169

Abstract: In this thesis I relate two theories about the content of phenomenal mental states. In plain language, these are theories about the way in which experience presents the world to the subject. Conceptualism is the theory that all phenomenal mental states are conceptually structured, so that their contents are limited by the conceptual repertoire of the subject. Nonconceptualism, on the other hand, is the theory that phenomenal mental states are not conceptually structured, and so their content is fixed in such a way that doesn't involve the subject's conceptual repertoire. I critique both theories and argue that neither conceptual nor nonconceptual content alone can account for all the features of phenomenal mental states. I conclude that a theory which allows for both conceptual and nonconceptual content is therefore desirable, and that we can avoid the obvious problem of attributing contradictory contents to phenomenal mental states if we adopt the view that phenomenal mental states can be stratified. In the final chapter I present some empirical evidence for a stratified view of phenomenal states, and give some indication of how to further proceed with a positive account of a mixed content theory.

Acknowledgements

I would like to acknowledge my committee: my supervisor Dr. Brian Garrett; my second reader Dr. Brigitte Sassen; and my third reader Dr. Mark Vorobej. Thank you all for your patience with my false starts, as well as for stimulating discussions and helpful comments on the paper. I owe an intellectual debt to Dr. Murray Miles who helped me early on in my academic career in getting my head screwed on straight, philosophically speaking. I would also like to acknowledge Dr. Gary Madison, with whom I've had countless fruitful discussions which have informed this thesis.

I was lucky enough to spend these two years with a group of fellow grads whose company I enjoyed and whose philosophical opinions I respected. My thanks to them, and best wishes for the future! Thanks to Sarah, with whom I have happily shared every moment of my philosophical journey. And to my family: my mother, father, and my sister Brandy, who always supported and encouraged me, and without whom I would never have completed this paper, or even had the opportunity to write it.

In Memory of Wayne Basil Pelette

"Time is but the stream I go a-fishing in."
- Henry David Thoreau

Table of Contents

1. Overview and preliminaries

"Nothing is more difficult than to know precisely *what we see.*"

- Maurice Merleau-Ponty[1]

1.1 Introduction

This paper is an attempt to characterize the content of those aspects of mental life that have a conscious – or phenomenal – component. As we will see, this topic has given rise to a vibrant, longstanding debate within the philosophy of mind. It is fair to say that the debate over the content of phenomenal mental states has largely broken into two camps: one which holds that all phenomenal mental states are conceptualized, and the other which believes that all phenomenal mental states are nonconceptual.[2] Following the favoured terminology, I will refer to these views in this paper, respectively, as "conceptualism" and

---

[1] Merleau-Ponty (1962) p. 67

[2] This opposed pair is the most widely accepted means of drawing the distinction. However, there have been other suggestions concerning how the distinction should be drawn. For instance, in his (1981), Fred Dretske draws the distinction in terms of analog versus digital content. In his (1995) he reconsiders, and redraws the distinction as acquired versus systemic content. For a discussion which contrasts these three positions, see Carruthers' (2000), specifically the fifth chapter. In the final section of this paper, I will consider Carruthers' claim that adopting the analog/digital distinction allows us to make sense of a mixed view of content.

"nonconceptualism."

For now I will give only a rough approximation of these views. As we will see, the debate presupposes that mental states – both conscious and non-conscious -- are representational in nature. Mental states represent the world as being some way or other, and a curious subset of mental states represent the world in such a way that the representation is characterized by a "look," or a "qualitative" aspect. The guiding question in this literature can therefore be stated roughly as follows: "What is the content of these phenomenal mental states?"

The conceptualist's view is that phenomenal mental states contain entirely conceptualized (or conceptualizable) content. There are several different ways in which to cash out the conceptualist's thesis. For instance, there is the claim that what we immediately see, what goes into making up the "look" of the perceived world, is a conceptualized environment of recognizable objects: people, places, things, and so on. Another would be to say that our phenomenal perceptual states are just the acquiring of (or, alternately, dispositions for acquiring) beliefs, and so there can be no act of perception which does not (or could not) give rise to some belief or other in us. If all experience is the acquiring of beliefs then the content of experience must be such that beliefs can "take hold" in them: having beliefs, then, presupposes the possession of concepts. However this position is cashed out, all conceptualists agree that if all mental states bearing conceptual content

were absent then we would see nothing. More evocatively, we might "see" swirling undifferentiated colour sensations, "hear" unintelligible sounds, and so on, but there would be no *sense* to it, no way in which to tell one perceptual event from another (or perhaps even the deliverances of one sense modality from another). The world would be like that of William James' infant: a big booming buzzing confusion. But a mere state of receptivity to the world does not amount to representation. It is conceptualized content, then, which fixes the look of our conscious perceptual states. Concepts are *in* our experiences, *inhabiting* them omnipresently.

A proposed (and polar) alternative to the conceptual view of phenomenal mental states is non-conceptualism, which most often identifies our phenomenology with a *sensitivity* to the environment, both distal and internal, in which no concepts are involved. On this view, cognitive activity invoking concepts is performed *upon* phenomenal states which are themselves entirely non-conceptual. Again, in the case of perception, what is immediately given to us are non-conceptual perceptual states made ready to enter the conceptual system as input, an array presented to our cognitive mind via our sensitivity to the environment. The look of perception is therefore put together *prior* to the application of concepts. Central to these views is that conceptualization is a *cognitive* task, and so should be conceived separately from the *informational* task

of representing the environment. Though these representations are accomplished by sub-personal means, they are nonetheless *presented to* the cognitive mind, as opposed to being *put together by* it.

The critical task which will make up the bulk of this paper consists of two parts. In **Chapter 2** of this paper I will present a critique of the conceptualist's position, which will mostly focus on the influential theory developed by John McDowell in his book Mind and World. I do not wish deny that cognition taking the form of conceptualization plays a *vital* role in understanding phenomenal mental states, but only that the role it plays must be properly circumscribed. If we think a state has phenomenology only insofar as it overtly or tacitly employs concepts via sub-personal cognitive routines or dispositions to engage in active cognition – such as in the formation of beliefs and judgements – we overlook the distinctive contribution *sense experience* makes within perception.[3] In such a view no place would be afforded to sensation in our conscious dealings with the world and I think that this is an undesirable conclusion.

In **Chapter 3**, my critical task will be to examine the arguments for non-conceptualism. The subject of most of my attention in that chapter will be Fred Dretske's theory of phenomenal mental states. My position will be that, though it

[3] Dretske (1981) p. 135

seems dubious in the extreme to claim, as conceptualists do, that there is no aspect of *sensation* within our conscious perceptions, to go to the opposite extreme and gloss phenomenology with sensation creates its own intractable problems concerning how cognitive states relate to experience. And much like the conceptualist's position, nonconceptualism leaves us with an unsatisfying description of our phenomenology. This is at least what I hope to show.

These theories each fail for their own reasons, and yet I would like to suggest that they both run into serious problems when we actually turn to our *phenomenology* in introspection and assess how things *seem to us* from our subjective point of view. I believe that neither does enough justice to experience, though each claims that their theory renders a more accurate picture of experience than the other. We might be tempted to think "So much the worse for experience!" But if any problem in philosophy has to answer to the way in which things appear to us, then surely this one must. To see this we need only consider that many of the arguments for either side take the form of observations from normal perception. Clearly, the question of content is tied down in many places to the methods of direct first-person observation.[4]

---

[4] Ray Jackendoff defends something like this view in his book Consciousness and the Computational Mind. In his view, if we believe that the phenomenological mind is a small subset of the "computational" mind, we can also reasonably claim that a phenomenological distinction will be caused by,

These two theories are mutually exclusive: a conceptualist will deny that sensations are experiences at all, and so even if they are content-bearing states it doesn't matter for the question at hand just what sort of content they have; on the other hand, a nonconceptualist will deny that concepts determine content, since conceptualizations are not phenomenal states, but only cognitive states bereft of any "feel." But conceptualism and nonconceptualism make up the opposing poles of the debate; there is, of course, a more prosaic, intermediate position which suggests that phenomenal mental states at times have both conceptual and non-conceptual content. Call this position the "mixed content" thesis. The thesis of this paper is that just such an intermediate or mixed view on the content of phenomenal mental states is most plausible. The essence of this view is nicely encapsulated by Peter Carruthers: while experience is *imbued* with concepts, it is not *exhausted* by them.[5] In **Chapter 4** of this paper, I hope to make a plausible case for the middle ground between these two positions. I will also consider the impact a mixed content theory of perceptual content might have for theories of

---

supported by, or projected from a computational distinction. What follows from this is that the "computational theory must be sufficiently expressive (must contain sufficient distinctions of the proper sorts) to make the world of awareness possible. Thus, if there is a phenomenological distinction that is not yet expressed by our computational theory, the theory must be enriched or revised." [Jackendoff (1987), p. 25]

[5] Carruthers (2000) p. 130

perception generally. The shift from either/or to both/and suggested by mixed content theories is not without its price; indeed, the problems associated with ascribing two different types of content to one and the same state has been a factor in pushing theorists to either conceptualism or nonconceptualism. I agree that there are difficult issues surrounding dual content, but they are wholly terminological, merely a product of how the debate is typically structured. As we will see, the moderate position may force a complete reworking of the distinction which the conceptual/non-conceptual pairing was meant to capture in the first place, and it may have even more radical consequences for the representational theory of mind in general.

But first, here in the remainder of **Chapter 1**, I will set about laying the groundwork for understanding the debate in general. I believe the debate as a whole can best be understood as the result of a confluence of two philosophical streams. One was a development within the philosophical literature on *consciousness* in which phenomenal consciousness – that is, the qualitative aspect of experience – was set apart from other aspects of consciousness as a problem in its own right. The other development was in the philosophical (and psychological) literature on *perception* which is concerned with properly

characterizing the type of content involved in perceptual states.[6] As a beginning I will briefly introduce these two philosophical streams. First I will explain what is meant by phenomenology, and show how it creates problems for naturalism. Then I will go on to outline one research program, representationalism, which aims at naturalizing phenomenology by reducing it to representational states with a certain content or role in the functional economy of the mind.

1.2 Phenomenology

A great deal of this debate centers around *perceptual* mental states, since they have the most *overt* phenomenology. Among perceptual mental states, those produced by vision are most often discussed. It is necessary right at the outset to give some indication of what I mean by phenomenal perception. Phenomenal perception is the way of perceiving that is possessed by putatively normal, language-using human adults *at the very least*. Whether or not similar perceptual

---

[6] The "phenomenal consciousness problem " antedates the "content problem" in the contemporary literature, with Thomas Nagel's 1974 article "What is it like to be a bat?" arguably serving as the first succinct presentation of the former, while Daniel Dennett's 1969 book Content and Consciousness is often considered the starting point for contemporary discussions on mental content. However, the problem of perceptual content has been around in one form or another since representationalist theories of mind first gained popularity beginning with Descartes, and so this problem antedates discussions of phenomenal consciousness in the wider philosophical canon.

abilities are possessed by pre-linguistic infants or higher animals is certainly an open question, and the views considered here point toward several possible answers. As "perception," phenomenal perception depends on our being sensitive to the environment through our sensory organs. But sensitivity falls far short of being phenomenal perception. We are not merely "sensitive" to our surroundings when awake, we are aware of it in a qualitative way.

A nice little piece of well-worn jargon which is helpful in sorting out the meaning of "phenomenal perception" is *quale,* or *qualia* in the plural. A quale is a qualitative or phenomenal property inhering in a sensory state. A quale could be the colour of an after-image, or that of an actual shade of orange; the various permutations in timbre and pitch in music; the smell of pine needles; the taste of an apple; the felt texture of silk. What we are conscious of in our perception are many different *qualia*-laden states: it is "like something" to undergo a phenomenal perceptual state, and what it is to see a tree is *like something* quite different than what it is to see a car or to smell a flower.[7] Mental states are *like something* because they are qualia-laden.

---

[7] Here I follow Nagel's apt and influential phrasing. From his (1979) collection, p. 166, here is a relevant excerpt from his canonical article "What is it like to be a bat?": "But no matter how the form [of consciousness] may vary, the fact that an organism has conscious experience *at all* means, basically, that there is something it is like to *be* that organism."

Most often, phenomenal perception is introduced by way of such examples which appeal to the subject to check his own experience, since at least in the case of language using humans we can recognize phenomenal states in ourselves in such a way that does not require any inference; they are immediately presented to us as feels or qualities. So we can say that a perception is phenomenal if it *appears* in some way or other to you. "Appearance," in this sense, is synonymous with "phenomenal percept."

Beyond these basic statements, it is no easy task to say further what is meant by phenomenal perception. As it turns out, telling a satisfactory story about what phenomenal perception is and how we receive it is more complicated than simply opening one's eyes. But one useful way in which to approach a further characterization is through contrast with non-conscious – or non-phenomenal – perception. It is virtually beyond question that perception is in part, if not *largely*, a non-conscious affair. It will be useful as a beginning to briefly defend this claim, so as to bring the puzzle of conscious perception more keenly into focus.

Prima facie, the very idea of a state of non-conscious perception is an oxymoron. Conversely, this is as much to say that the phrase "conscious perception" is redundant; seeing something is simply having it presented *to you*, and *ipso facto* this means perception must always be visually *like something or*

other for the perceiver to undergo.[8] For example, on this account, a mental state can only count as a perception of the moon if the moon *appears,* in some way or other, in the conscious visual field. A mental state could not be a perception of the moon if that state wasn't *like something or other* for the perceiver to undergo, in which case it could no more be *of* the moon than *of* a giraffe. It follows from this that a phenomenal perceptual state is characterized, or individuated, with reference to the "look" it manifests, as opposed to, say, the functional role it plays: same "look," same perception. Though this is an incomplete description of perception in general, it adequately describes what is meant by "phenomenal perception."

That some measure of phenomenology is intrinsic to perception seems like a fair conclusion to draw when judging from one's normal experiences, since few

---

[8] This idea holds a certain appeal to common sense, and yet it is difficult to find a philosophical perspective on consciousness which explicitly endorses it. Descartes comes as close as anyone. Consider this passage from the Fourth Set of Objections. Arnauld, who authored the objections, suggests that there can be things in our mind of which we are not aware, citing the example of a infant: "The mind of an infant in its mother's womb has the power of thought, but it is not aware of it." Descartes reply: "As to the fact that there can be nothing in the mind, in so far as it is a thinking thing, of which it is not aware, this seems to me to be self-evident. For there is nothing that we can understand to be in the mind, regarded in this way, that is not a thought or dependent on a thought." The case of the infant, Descartes holds, is simply a matter of our having forgotten what went on in our minds during our time in the womb, and not of our having had something in our mind of which we were unaware. [Quotations taken from Descartes (1641) p. 74]

perceptions occurring within the normal course of experience could serve as a counterexample. Despite this, examples of non-phenomenal perception abound, both in normal and abnormal instances of perception. To take an everyday example, as you sleep your body periodically adjusts itself in response to stressed postural positions. This is certainly an instance of the body being sensitive to the environment although, being unconscious, it is not like anything at all for you when your body makes these adjustments. There are also examples of non-phenomenal perception becoming conscious, though these are controversial. Take the example of a clock striking five. It is possible for you to miss the first three chimes in your experience, and yet, when you become aware that a clock is chiming – say on the fourth chime – you may be able to recall the occurrences of the earlier chimes and piece together how many have occurred.[9]

But the well-documented pathology commonly referred to as blindsight is perhaps the most compelling case of non-phenomenal perception or experience. Blindsight occurs in individuals who suffer from cortical blindness, which can be acquired congenitally or caused by trauma to the primary visual area (V1) of the

[9] I say that this is controversial since this interpretation of such non-conscious experiences has been challenged by Dennett (1991, p. 137-138), who sees these sorts of cases as our experiencing rapid, short term, "rolling" memory loss. In other words, we heard the chimes (ie. they were part of our phenomenology) but we almost immediately forgot about them until recalling them became a cognitively salient thing to do.

brain. To greatly oversimplify, the effect of cortical blindness is that it prevents the *conscious* processing of perceptual information. On the other hand, blindness caused by damage to the eye itself prevents information from entering the visual system *in toto*. Despite this physiological difference, the *phenomenological* component of cortical blindness – what it is like for the afflicted person – is identical to instances of blindness caused by trauma directly to the eye. This is, of course, why it is called *blind*sight. In both cases, the stricken individual insists that they have no experience of any sort in the area of their visual field affected by their condition. But unlike individuals who are blind in the common sense of the word, individuals with cortical lesions resulting in blindness often exhibit the remarkable, seemingly magical ability to make accurate perceptual discriminations in the affected areas of their visual field. For example, when (and *only* when) asked to indicate the position of a light source or the orientation of a line in their affected field they provide correct answers with far greater frequency than could be chance, even though they claim that each answer they provide is nothing more than a guess.[10] This result suggests that there is some manner of perceptual information being processed in the brain of the individual, strong enough even to guide her behaviour, without her having any qualitative apprehension of the

[10] See Weiskrantz (1974) for a detailed (and pioneering) study of blindsight patients.

perceptual event in which the discrimination is actually made. Though consciously blind, she is non-consciously sighted.[11]

So perception is not only a matter of the world appearing in some way or other to the observer. Perception also occurs in creatures out of which, some contend, no *subject* or *observer* is peering, as well as in sub-personal "parts" of creatures which *do* consciously enjoy phenomenal perception, such as ourselves. This will be a crucial distinction in the following discussion: it is one thing for *you* to enjoy a perception, and it is quite another for only some part of the brain to have access to perceptual information. Just where to draw that line is the question at hand: At what point does the phenomenal interface between organism and world (rather than brain and sensation) become implicated in perception?

1.3 Naturalizing phenomenology

Beyond simply finding a definition of phenomenal perception and distinguishing it from other forms of perception, there is another side to this literature, equally important for our discussion, which suggests that the qualitative aspects of experience create intractable problems for a naturalistic view of the

---

[11] I am not suggesting that some sorts of non-phenomenal perceptual content cannot come to be phenomenal, just that not all perception is, at all times, phenomenally experienced.

world; that qualia cannot be given an adequate scientific explanation.[12] Colin McGinn regards the enterprise of naturalizing phenomenal consciousness as doomed from the outset. For McGinn, phenomenal consciousness will remain a mystery because our minds are not equipped to understand such things. This is not to say that there *is* no natural explanation of phenomenal consciousness, but rather that we have no access to such an explanation.[13] As in the case of conscious vs. unconscious perception discussed above, we can better understand theories which attempt this reduction (such as the theories under consideration here) by setting them off against these so-called "Mysterian" positions which hold that phenomenal mental states will always remain some sort of mystery.

McGinn's pessimism stems from his belief that a satisfactory explanation of consciousness is "cognitively closed" to us; our minds simply lack the capacity

---

[12] We need not be concerned much here with exactly what "naturalization" would entail. But whatever naturalization might mean precisely, it would involve turning the "mystery" of phenomenal consciousness into an approachable scientific problem by taking it apart and asking pointed questions about its structure and function. "Naturalization" is often mentioned in the same breath as "physicalism" or "materialism," since, on all accounts, naturalization proceeds by showing how a phenomenon can be explained entirely in physical or materialistic terms.

[13] McGinn's argument is not meant to establish the falsity of naturalism, that is, that all natural phenomenon could in principle be adequately explained. But it is a challenge to the naturalist's project, and therefore, to the theses considered in this paper which assume that an understanding of phenomenal consciousness is within our grasp.

to comprehend a naturalistic account of phenomenal consciousness. The justification for "cognitive closure," in turn, is based in part on Jerry Fodor's thesis that the mind is largely modular in its structure. According to Fodor, our minds, having been cobbled together through evolution, have cognitive limitations. They are constructed from a variety of innate and specialized modules which are each attuned to a certain task, such as language acquisition, folk psychology, and so on. These are modules insofar as they operate largely in isolation from other such modules, and so each module draws on its own distinctive sort of input and produces its own sort of output. The upshot of this for McGinn's thesis is that the epistemic possibilities open to any creature will be fixed or limited by the sorts of modules present and operating in their brains, since it implies that certain epistemic domains will be open to some creatures and yet "cognitively closed" to others.[14] For instance, the mind of a crab will simply lack the capabilities for understanding in certain domains where the mind of a cat excels. Or a man blind from birth will not be able to know what sighted people experience when seeing colour. As for human beings trying to grasp how subjective experience arises from inert matter, McGinn believes the enterprise to be about as promising as slugs trying to do Freudian psychoanalysis.

[14] McGinn (1990) p. 3

A similar point is made by Thomas Nagel (also cited by McGinn as an influence on his idea of "cognitive closure") who believes that an understanding of the nature of phenomenal consciousness enjoyed by other creatures is impossible for us due to the particularities of our senses. For instance, he writes:

> **. . . in contemplating [the phenomenal experiences of] bats we are in much the same position that intelligent bats or Martians would occupy if they tried to form a conception of what it was like to be us. The structure of their own minds might make it impossible for them to succeed.**[15]

But, Nagel contends, just because the Martians or sentient bats cannot understand what our experiences are like, it does not follow that there are no *facts* about the nature of our experiences. Obviously there *are* such facts since *we* know about them first-hand! The Martians and sentient bats are simply "cognitively closed" to them, just as we are to the facts about their experiences. Furthermore, if we tried to discover what it is like to be a Martian or a bat through scientific enquiry we would fare no better. In seeking objectivity we try to find an explanation of the world that does not depend on the peculiarities of an individual's sense organs or limited perspective. But since science purports to be a "view from nowhere," it cannot therefore hope to account for a view from *somewhere*. So, according to Nagel, there are limitations to the objective view of the world: " . . . there may be aspects of reality beyond [human objectivity's] reach because they are altogether

---

[15] Nagel (1974) p. 440

beyond our capacities to form conceptions of the world."[16]

I do not think the notion that an explanation of phenomenal consciousness is cognitively closed to us stands up to scrutiny. To begin with, those who endorse cognitive closure seem to argue from the fact that our minds contain specialized modules which are meant to operate in certain domains to the assertion that there are particular domains, presumably domains for which we have no specialized modules, which are cognitively closed to us. This is unwarranted; it may simply be that domains for which we have no specific knack are more difficult for us to make progress in, but not impossible. We may have to draw together resources from several different modules in order to make slow progress in these areas, or perhaps employ some less finessed general purpose learning mechanisms.[17] But McGinn does not consider these possibilities.

Further, McGinn glosses over an important difference between the slug's relation to Freudian psychoanalysis and our relation to the naturalization of phenomenal consciousness: whereas as we can express incredulity in the face of questions which phenomenal consciousness poses to us, the slug cannot even begin to comprehend the questions inherent in psychoanalysis. Unlike the slug,

---

[16] Nagel (1986) p. 91

[17] This point is made by Carruthers (1996) p. 61

we understand the unanswered questions regarding consciousness. But it would be strange indeed if we were able to grasp these sorts of questions and yet have no way of approaching them. They are not, after all, questions on the order of "How many angels can dance on the head of a pin?" or "What does Thursday weigh?"

Beyond arguing for cognitive closure, McGinn has another argument for believing that an understanding of phenomenal consciousness lies outside of our ken. McGinn suggests that there is an "explanatory gap" between the qualitative aspects of experience and any sort of physicalist or materialist explanation which would try to account for those qualities. It should come as no surprise that most philosophers fond of pointing to this gap are pessimistic about its ever being closed, or obstinate in their belief that it cannot be closed due either to the nature of phenomenal consciousness or to our own epistemic failings. In McGinn's case, his pessimism stems from the latter. If we are going to attempt to naturalize phenomenal consciousness and understand how chunks of matter take on an inner life, says McGinn, we really only have two tools at our disposal: the first-person point of view we enjoy in introspection and the third-person point of view taken by science. Either we use the tools of introspection to "dig deeper" into our experiences in an attempt to develop a more sophisticated set of phenomenal concepts which might get us closer to a naturalistic explanation. Or we work from the other end by investigating the physical structure of the brain, hoping that an

explanation of phenomenal consciousness will miraculously turn up. But McGinn believes that neither of these approaches is pitched at the right level; we cannot understand how consciousness arises in matter by closely scrutinizing either one or the other of consciousness or matter.

This much is certainly true. No matter how hard we introspect we won't ever come to understand how consciousness arises in neurological events. As Dennett points out, "we have no direct personal access to the *structure* of contentful events within us."[18] That is to say, in introspection we are presented with certain content-bearing states, but we are not thereby given any information as to how this content is accomplished (Dennett assumes, plausibly, that this is a function of the state's structure). It's hard to resist the analogy here with a computer interface, which we deal with directly, and the unseen software which makes the interface possible, so I won't.

And on the other hand, no matter how closely we study neurological events, we will never understand how they give rise to phenomenal consciousness. The structure and activities of a neuron do not seem to suggest that they give rise to mentality. If a creature who had no notion of phenomenality were to study both brain cells and heart cells side by side he/she/it would never be able to discern that

[18] Dennett (1978) p. 103

one type of cell gave rise to qualitative states while the other did not. There would be no empirical reason for the creature to postulate phenomenality to be an effect of certain configurations of neurons.

What we would need, McGinn suggests, is a sort of middle domain from which to work, one which did not depend for its conceptual repertoire on either the domain above (phenomenal consciousness) or below it (neurophysiology). In this middle domain, McGinn suggests:

> **. . . operative properties would be neither at the phenomenal surface nor right down there with the physical hardware; they would be genuinely deep and yet they would not simply coincide with physical properties of the brain.**[19]

Unfortunately for us, it is just this sort of domain the exploration of which McGinn thinks is cognitively closed to us, for, as he says, its "characterization would call for radical conceptual innovation (which I have argued is probably beyond us)."[20] We might say that a solution to the problem of phenomenal consciousness falls through the holes in our epistemological net.

Several of McGinn's critics are baffled by his bold-faced claim that there are no cognitively accessible intermediate domains between phenomenology and neurology, and rightly so, since there are many different levels of scientific

---

[19] McGinn (1990) p. 103

[20] Ibid. p. 102

enquiry between them.[21] Among these intermediaries are "a variety of forms of computationalism, together with kinds of functional description characteristic of much cognitive psychology."[22] So it is no wonder that McGinn would feel that it was a mystery as to how consciousness arose from brain states, since he has chosen to disregard all of these intermediate scientific steps on the way to an explanation of consciousness.

But the most crucial consideration which McGinn overlooks is the possibility that there is more than one way to proceed "downward" from phenomenal consciousness. The only method he considers is introspection, and, as we have seen, this method is out of the question; as helpful as it might prove to be in solving certain problems, it doesn't get us anywhere with this particular one. What McGinn does not consider is that we might close the gap by re-conceiving phenomenal consciousness in something like intentional or functional terms, or, put differently, that we should ask ourselves what sort of cognitive state it would

---

[21] In fact, in his Times Literary Supplement review of McGinn's book The Problem of Consciousness (in which he sets out the "explanatory gap" argument), Dennett cuttingly points out that the desiderata McGinn suggests for an intermediate level theory – that is, a radical conceptual innovation autonomous from the conceptual framework of the adjacent projects – is in fact an excellent description of "exactly the set of concepts that are being developed in cognitive science ... but also increasingly exploited in cognitive psychology and computational neuroscience."

[22] Carruthers (2002) p. 63

make sense for phenomenal consciousness to be. (Curiously, he seems to consider the analogous position working upward from brain states; he talks about an attempt to discern the phenomenal within brain states and is, of course, pessimistic about the prospects of this approach.) In other words, we attempt to reconstruct the level *immediately* below consciousness. Once we find an acceptable cognitive state which captures consciousness in a non-circular way – that is, a cognitive state that is not defined by way of the terms related to consciousness – then the prospects of naturalizing phenomenology will start to look a lot more promising.[23] The representational theory of mind, which will be

---

[23] There is the possibility of a *major* tangent here (that I will resist, but for this footnote) regarding David Chalmers' version of the explanatory gap, which remains as wide as ever despite attempts to offer functional or intentional explanations of phenomenal consciousness. For Chalmers, phenomenal consciousness is, by its nature, not the sort of thing that can be given an explanation in physicalist terms. Chalmers' obstinate refrain in the face of these theories (indeed, in the face of any reductionist theory of consciousness) is that we can coherently imagine (ie/ it is a logical possibility) that the functional/intentional states that are said to be identical with phenomenal consciousness could accomplish their work "in the dark,", that is, without manifesting in a phenomenally conscious way. For Chalmers, these theories cannot explain why, for instance, states playing a certain functional role come to be conscious, but only insist that they do. Another way in which Chalmers is fond of making this point is to say that these theories offer "correlations" in place of "explanations."

There are at least two ways to attack Chalmers' position. One way is to argue that "conceivability" is not the best guide in deciding what is logically possible. For instance, it is certainly *conceivable* that Batman and Bruce Wayne are distinct individuals. But I would simply be mistaken if, having conceived of them as different, I asserted that it was logically possible for Bruce Wayne to be at Wayne Enterprises and for Batman to be in the Bat Cave at the same time. It is

examined in this paper, adopts exactly this strategy in its approach to phenomenal consciousness.

1.4 Representationalism

The preceding discussion of phenomenal consciousness has provided us with a rough characterization of what is meant by the term as well as an introduction to how we might proceed with its naturalization. The positions considered in this paper can best be understood as attempts to close the explanatory gap separating phenomenal states and neurology using the top-down approach discussed directly above. To this end, both conceptualists and nonconceptualists agree that phenomenal states can be understood as states

---

not logically possible for this to be the case, since Wayne and Batman are *in fact* the same individual. Applying this to Chalmers argument: simply because we can conceive of all the psychological aspects of consciousness being in place while no phenomenal aspect of consciousness manifests, it does not follow that this is a logical possibility. Conceivability therefore is not a reliable guide to logical possibility.

Another way to attack Chalmers is to explain the explanatory gap itself, as Chalmers understands it, as some manner of cognitive illusion (though the wider, yet more approachable, gap McGinn speaks of is still very real). After all, we don't bother raising similar protests about reductive explanations of digestion by suggesting that it is logically possible to imagine the action of enzymes in the alimentary canals of human beings going on without any digestion being involved. [Tye (2000) p. 22] If we were to raise this sort of protest, we would simply be accused of misunderstanding the word "digestion." I am suggesting that the case of phenomenal consciousness is analogous to this one.

bearing a certain form of representational content; they simply differ on the fine point of what sort of content this will turn out to be. But before we turn to a consideration of the theories themselves, I would like to say a little more about the family to which both theories belong, that is, representational theories of mind. I will also discusses the genus "content," to which both conceptual content and nonconceptual content belong.

Michael Tye begins his overview of phenomenal state representationalism by stating that it "is a thesis about the phenomenal character of experiences, about their immediate subjective 'feel'."[24] This characterization is fine for our purposes, but it only tells half of the story. Representationalism is, in fact, a theory of both the phenomenal character of experience *and* of the character which belongs to non-phenomenal mental states; in its strongest form, it is meant to account for both. In what follows I will attempt to answer the question "What does it mean to say that a mental state is representational?" by picking out some key features of mental representations.

A. Representations fall between stimulus and behaviour. The need to appeal to representations first arises (and the explanatory gap first opens) in instances of sensation where, as Jose Luis Bermudez puts it, "the law-governed

---

[24] Tye (2000) p. 45

correlation between stimulus and response breaks down."[25] There is no need, for instance, to appeal to representations to explain the behaviour of wood ticks, which possess only a limited repertoire of behavioural reflexes in response to simple environmental triggers. Ticks will find a high perch and remain there indefinitely, or until the smell of butyric acid, secreted by the skin of all mammals, triggers them to loosen their grip and drop down from their perch. Landing on some solid surface extinguishes the initial behaviour and serves as a cue to initiate a search for heat. When the tick locates a warm spot, this serves as a third environmental trigger which overrides the second trigger and its "searching" behaviour, and the tick begins burrowing down toward the heat. Each response is law-governed and doesn't allow for any cognitive flexibility.[26] Jacob Von Uexkull – who wrote imaginatively about the perceptual worlds (*Umwelt)* of various creatures – evocatively describes the ticks 'world' as follows:

> **The whole rich world around the tick shrinks and changes into a scanty framework consisting, in essence, of three receptor cues and three effector cues – her *Umwelt.* But the very poverty of this world guarantees the unfailing certainty of her actions, and security is more important than wealth.**[27]

The behaviour of the tick is like a mousetrap being sprung when it "detects"

---

[25] Bermudez (1998) p. 88

[26] Carruthers (2000) p. 124

[27] Von Uexkull (1934) p. 12

pressure. In explaining this behaviour we do not need to appeal to an intermediate state between the stimuli and the response where the environment is represented to the tick, since the tick's response is invariant. But while an explanation along these lines will suffice for the tick, it will not suffice in our own case. Our minds are not mousetraps waiting to be sprung by certain stimuli. There is an element of *representation* in our perceptions, which puts a distance between stimulus and response. Instead of our behaviour merely being triggered by direct environmental pressures, the world is presented to us as being some way or other so that we may act upon it in a more flexible and dynamic fashion, in ways which are unavailable to a wood tick.

B. Representations are about worldly properties. According to representationalism, what is represented in a mental state is a grouping of features inhering in the world, and not some purely mental phenomenon or sense-data. The phenomenal character of an experience is therefore due to its representation of objective, non-experiential properties. In fact, one of the great strengths of the theory is that it can provide an explanation for why it seems to us that we perceive parts of the world outside us *beyond* our own mental states. This feature of perception is often referred to as "transparency" or "diaphanousness." A quick contrast between sense data theorists and representationalists will help to show the importance of this feature of perception.

Sense data theorists and their ken are committed to the view that the transparency of perception is a massive cognitive illusion since, on their account, we don't see the qualities of the objects themselves but only the qualities or properties of our own experiences; the outside world is only "seen" through seeing some intermediary sense-data. Similarly, Mysterians argue that the phenomenal aspects of our perception are something over and beyond the representational aspects, so that it is at least logically possible that they should come apart. But why should we take our experiences to be *about* the world, and why should we take them to *matter* to our representing the world? It seems the only answer that can be given here is that perception systematically misleads us. This conclusion is certainly problematic, and should motivate us to look for an explanation which would not regard our belief that we directly see the world as a ruse.

How does the representationalist handle the problem? Normally we see *through* our perceptual states and hardly take any notice of them at all *as* experiences; our attention is instead on the objects in the world with which we are dealing. Even when we attempt to turn our attention away from our worldly dealings and examine the experience itself, we are again and again lead back to the things in the world. You are walking through the woods and you stop to look at particularly nice-looking tree. As often happens when looking at trees, a contemplative mood descends upon you, and, as a philosophical exercise, you

attempt to see *beyond* the tree. That is, you attempt visually to *get past* the tree that is in front of you and through to the intrinsic qualities of the experience itself, something other than what the experience is *of*. Is this possible? According to Michael Tye, it isn't: "one's awareness seems always to slip through the experience [to the tree-ness and the green-ness], as instantiated together in an external object.[28]" That experience appears to be transparent is important because it makes a strong intuitive case for there being no phenomenal features – properties of experiences – over and above representational features. Therefore:

> **. . . the phenomenal character of such experiences – itself something that is introspectively accessible, assuming the appropriate concepts are possessed and there is no cognitive malfunction – is identical with, or contained within, their intentional [or, what is the same, representational] content.[29]**

C. Representations are contentful. In A and B, I specified a functional and, we might say, a teleological "mark" of representations, respectively. But the most important feature of representations for our purposes is that they are content-bearing states. Representational – or intentional – states, as we have seen, are about the world; the content of a representational state is simply *what* it is about. Another way of saying this is that the term "content" refers to the way in which "some aspect of the world is presented to a subject; the way in which an object or

---

[28] Tye (1996), p. 30.

[29] Ibid. p. 136

property or state of affairs is given in, or presented to, experience or thought."[30] For instance, my wishing for a short winter is a mental state characterized by the attitude of wishing, and with the content "the winter will be short." I can also have representational states which are not about any actually existing entity. My belief that chupacabras exist represents the world in a certain way, albeit a way that is false. Nevertheless, both my hope that winter will be short and my belief in chupacabras represent the world in a certain way.[31]

Representations also prescribe what the world would have to be like for that representation – my hope, my belief – to turn out to be either true or false; they are governed by what theorists call *semantic normativity*:

> **. . . whether its content is conceptual or nonconceptual . . . an intentional state presents the world as being a certain way; and intrinsic to this presentation, to its content, is a set of (semantic) conditions under which it does this correctly, truthfully, satisfactorily . . . and so on.**[32]

So my belief that chupacabras exist represents the world as containing small, blood-sucking, half-alien, half-dinosaur, tailless vampires with quills running down their backs, and as also involving a massive conspiracy initiated at the

---

[30] Cussins (1990), p. 133

[31] This belief need not be occurrent in order to be a representation. I can be said to believe in chupacabras whether or not I am currently conscious of an occurrence of this belief. So non-conscious mental states can also be representations.

[32] Gunther (2003) p. 5-6

highest governmental levels to cover up these twisted abominations from an alien's laboratory (that is, if I were both a crypto-zoologist *and* a rabid conspiracy theorist, and *why not?*). If these conditions are not met – if these things do not really exist in the world – then my intentional state is a *mis*representation. It is therefore intrinsic to representations, as content-bearing states, that they can *mis*represent or that their content can be incorrect.

A, B, and C do not exhaust the characteristics of representational states. For instance, many theorists would want to extend representations so that they would include other forms of normativity, ie. *rational normativity*. They may be right to extend "representation" in this way, but this issue will not be important for the following discussion, while the features of representations I have picked out *will* matter for the discussion ahead. A and C especially will arise in the context of the following debate as important factors.

Now that we have a general picture of what representationalism is, let's return to Tye's initial claim that representationalism can be used to explain phenomenology. While some theorists will want to suggest that being a representation is *all* that is required for a mental state to have phenomenology, most believe that a state's being representational is *necessary* for it being

phenomenal, but not sufficient.[33] It must, in addition, meet certain other criteria. So a phenomenal mental state must then be a representational state, "**A**", plus some "**x**" which, loosely speaking, *makes* the representational state a conscious one. Now, "**x**" could turn out to be some other sort of mental state which interacts in the right way with "**A**", but it could also specify a *place* in the functional economy which "**A**" must occupy in order to have phenomenality conferred upon it, or something quite different altogether. But regardless of these differences, "**x**", for a representationalist, will certainly not be anything ontologically novel.[34] It is this "**x**" which the debate we are about to enter into circles around.

1.5 Kinds of content

One of the major difficulties in assessing this debate are the countless ways in which the term conceptual content has been understood. Simply put, there does not seem to be any canonical way, agreed upon by all or even most conceptualists, in which to fix its extension. Some conceptualists have very weak

---

[33] But this would gloss over the distinction between conscious and non-conscious perception outlined above.

[34] Representationalists are therefore necessarily anti-qualia, at least in the sense of the term qualia in which it attached to Mysterian arguments, in which they are taken to be non-functional aspects of experience. Of course, representationalism should not be understood as denying that there is a real difference between conscious and non-conscious states.

conditions which content must meet in order to be considered conceptual, while others have a more narrow and demanding view of conceptual content. Nevertheless, there is a general framework developed by York Gunther in his introduction to the debate that can be applied to all theories which claim that experience has conceptual content.

In the course of giving an exacting picture of what is meant by "conceptual content" Gunther begins by invoking Frege and the distinction he made between "sense" and "reference," and suggests that the practical meaning of "conceptual content" is similar in nature to what Frege meant by "sense."[35] Frege in part developed his theory of "sense" in response to what is now, in light of his analysis, a well known philosophical puzzle about identity statements, which I will now recount.

Frege's theory is thought to correct certain shortcomings in John Stuart Mill's *reference* theory of meaning. Without going into details, the *reference* of an expression is "a function of the reference of its parts (the thing[s], propert[ies], relation[s], state of affairs, and so forth)[.]"[36] For instance, the statement "I have a

---

[35] This discussion of Frege's relevance to the debate is largely taken from York Gunther's "General Introduction" to the (2003) collection which he edited, p. 7-14.

[36] Gunther (2003) p. 7

red book on my wooden shelf," refers to certain things in the world (a book and shelf), with properties (being red, being made of wood), standing in a certain relation (of one being on the other). Now, the philosophical puzzle: if, as Mill believed, the meaning of a proper name was its *reference*, then where *a* and *b* refer to the same object "statements of the form $a = a$ *should* have the same content as statements of the form $a = b$."[37] And yet it seems obvious that there are statements which violate this rule. For instance, under Mill's theory of reference the identity statements "Doctor Jekyll is Doctor Jekyll" and "Doctor Jekyll is Mr. Hyde" would have precisely the same *meaning*, since both "Dr. Jekyll" and "Mr. Hyde" are names for one and the same individual. Both statements should therefore be *trivially true* given that they have the same reference and therefore, according to Mill, the same content.[38] But this seems obviously wrong when we consider the *practical* impact these two statements might have on someone's beliefs; whereas the first identity statement is a simple tautology which would come as no surprise to anyone, the second would be genuinely instructive to everyone in Robert Louis Stevenson's book (consider Dr. Utterson's possible reactions to either statement!) which suggests that the content of the second

[37] Ibid. p. 7

[38] Ibid. p. 7

statement diverges critically from the meaning of the first.[39]

In light of such examples, Frege suggested that a pair of proper names like "Dr. Jekyll" and "Mr. Hyde" must, *contra* Mill, express different contents; their meanings therefore cannot be fixed by their *reference*, which is identical, but instead by what Frege called their *sense*. An expression's *sense*, we might say, is the 'cognitive value' which it has for the person who possesses it; the sense of an expression can therefore differ from person to person even though the expression might *refer* in each case to the same entity. By introducing the notion of sense, Frege moves the content of an expression from its object to the 'understanding' of the subject who entertains it. Gunther therefore suggests that we can understand a sense as specifying a 'possible route' to a reference, so that "Dr. Jekyll" and "Mr. Hyde" specify unique routes to the same individual who is being referenced. Of course, as was established in the example above, for one to understand two senses of a term does not imply that one must also recognize that they are different routes to the same reference.

---

[39] If one is bothered by how the physical transformation Jekyll undergoes in becoming Hyde might skewer the example by introducing new properties, then we could resort to Frege's own stuffy example: "Hesperus" and "Phosphorus" are both names for the planet Venus considered in its appearance as the evening star and the morning star, respectively. It is clear that, though both names refer to the planet Venus, the identity statement "Hesperus is Phosphorus" might nonetheless be genuinely instructive to someone.

Though Frege developed this theory to account for features of language, his insights have naturally been extended to an account of mental states with normative conditions such as beliefs. Consider that Dr. Utterson, if asked where he believed Dr. Jekyll and Mr. Hyde were at the moment, might say something like this: "It is probable that Jekyll would be in his study at this hour. Hyde is likely lurking in some shadow, miserable creature that he is." Since on the Millian account "Dr. Jekyll" and "Mr. Hyde" would have identical content, we therefore would be forced to attribute two contradictory sets of beliefs to Dr. Utterson: "Dr. Jekyll is in his study" and "Dr. Jekyll is not in his study," as well as something like "Mr. Hyde is likely lurking in a shadow" and "Mr. Hyde is likely not lurking in a shadow." But the Millian analysis does not seem to capture the *sense* of Utterson's belief, since, being ignorant of the fact that Jekyll and Hyde are one and the same person, he *believes* that his statements concern two different individuals. On the other hand, in taking the Fregean line we could absolve Dr. Utterson of holding contradictory beliefs; the concepts which made up either of those beliefs – "Dr. Jekyll" and "Mr. Hyde" – though they referred to the same entity, would have distinct contents and therefore distinct meanings.

The benefits of adopting a Fregean framework for understanding mental states such as beliefs are obvious. Intuitively, Fregean senses seem to do a better job of capturing the viewpoint of the believer; it allows that believers can grasp

certain aspects under which a reference is thought while being oblivious to others. On a related note, adopting Fregean sense when interpreting another's beliefs accords with the Principle of Charity, which states that "by minimizing an individual's inconsistency one maximizes her intelligibility, arguably the point of interpretation."[40]

So the Fregean analysis offers a general framework for understanding what is meant by 'conceptual content.' Gunther then goes on to identify four principles of conceptual content which can be derived from adopting the Fregean framework; we will need only three for our purposes:

A. Compositionality
B. Cognitive Significance
C. Reference Determinacy

Each of these principles Gunther believes to be a necessary condition for 'conceptual' content. It will be useful to go through each of these principles, since each will figure in the discussion to follow.

A. The Principle of Compositionality states:

If content *c* is conceptual (and complex), then *c*'s constituents functionally determine *c*. [41]

---

[40] Ibid. p. 8

[41] Ibid. p. 9. Gunther mentions that his principle is nearly identical with Gareth Evans' Generality Constraint which a content must meet in order that it be conceptual. The Generality Constraint states:

To take a concrete example, the assertions "Pierre studies astronomy" and "Anna studies philosophy" both have their content in virtue of the concepts they each employ; in this case: *Pierre*, *Anna*, *studies*, *astronomy*, and *philosophy*. These concepts *are* the constituents of the content of either phrase. Complex conceptual contents are therefore "pieced together" from discrete components – concepts. This feature of conceptual content is devised to explain certain features of language, specifically its *systematic* and *productive* nature. In order to see how the Principle of Compositionality can do this, let us briefly examine the failure of the opposite position to do so.

One (not too popular) alternative to the view that propositional attitudes (like my belief that Pierre studies philosophy) are constructed out of discrete components knit through relations is that relations are semantically "fused" with their objects. This is what Jerry Fodor calls the "fusion story"; "fusion" in the sense that it takes every propositional attitude to be "atomic," that is, as being

---

For any two thoughts, *Fa* and *Gb*, and for any thinker *s*, if the contents of *Fa* and *Gb* are conceptual, then if *s* understands both *Fa* and *Gb*, *s* also understands *Fb* and *Ga*.

In other words, if you understand "Pierre studies astronomy" and "Anna studies philosophy," you must also understand "Pierre studies philosophy" and "Anna studies astronomy" in order to be said to possess the concepts "Pierre", "studies", "philosophy" and "astronomy." A failure to do this would mean that your knowledge of the meaning of, say, "philosophy" was context-bound, to the extent that you would be unable to extend its use to propositions involving "Anna."

made up of a referring expression ('Pierre') and a one-place predicate with no internal structure (which he renders as 'studies-astronomy' to express its atomism).[42] What this would mean is that propositions such as "Mary thinks-John-is-nice" and "Sam thinks-John-is-nasty" would "have no more in common at the level of canonical notation than say, "John eats" and "Mary swims.'"[43] Intuitively, this result would be enough to put the view to rest, but Fodor argues further that this makes it seem impossible that we could ever learn language in the first place, since on the fusion view there would be infinitely many sentences of the form "x believes-y" which, being atomistic, would have no relation to one another. The fusion view can neither account for the productivity of language, since in the propositions "John likes sailing" and "John likes walking" the reoccurrence of "likes" would be purely a matter of coincidence.

The view that we are considering, wherein each concept is generalizeable and able to be recombined in an infinite number of propositions, seems to better account for the way we learn language. Whereas the fusion view seems to imply

---

[42] Fodor (1981) p. 178. There are interesting affinities between the "fusion" view of propositional attitudes, which Fodor takes to be hopeless, and the story of nonconceptual content endorsed by Andy Clark in his (2003), in which he renders expressions of nonconceptual content with hyphens to show their context dependency.

[43] Ibid. p. 180

that language acquisition is accomplished serially, one whole phrase at a time, the compositionality of conceptual content instead suggests that having command of a concept implies that we can deploy it in novel situations as well as familiar ones.

B. The Principle of Cognitive Significance is intended to highlight the connection that conceptual content has with rationality. Gunther suggest that Cognitive Significance can be construed in two ways. On the weak construal, the content of a mental state is conceptual if it would lead the subject to have some beliefs corresponding to that state. The principle might be formulated as follows:

If the content c of a mental state *m* is conceptual, then *s* believes that c in normal conditions.[44]

So, if you are in a representational state with the content "I am in front of University Hall," then all things being equal you will come to *believe* that you are in front of University Hall. The contents of this belief are considered conceptual since they are "governed by rules of rationality"; we might say that the belief falls within the *space of reasons* – but more on this in awhile. Of course, the belief need not be given any explicit linguistic form, but might simply be reflected in your behaviour in light of the representational state; if you have a class in University Hall, you'll proceed up the front stairs and into the building, say. If, on the other hand, *s* did not come to believe *c*, then the content of *c* must be

---

[44] Gunther (2003) p. 10

nonconceptual, since it has "fallen through the net" of *c*'s rationality; it is belief-independent.[45]

Gunther believes that the weak construal is ultimately too vague. It just assumes that beliefs can always be given a rational grounding, and that lower-level mental states such as perception are not governed by rationality. But what is the specific notion of rationality being mobilized here, and why can't "lower" mental states be said to meet the conditions for rationality? The weak construal only seems to reiterate intuitions about what nonconceptual content might be without *grounding* those intuitions.

Gunther finds that the strong construal fares better in being explicit about its commitment to rationality. The principle states:

If content *c* is conceptual, then *F* and *G*, which are constituents of *c*, are different concepts if an individual could have an intentional state about a

---

[45] Of course, conceptualists will want to deny this possibility altogether. Since, for them, there is no such thing as nonconceptual content, there can be no state bearing a concept which does not give rise to some belief or other about that content in us. This will at least be true for conceptualists who adhere to something like the weak construal, such as David Armstrong. Although he put it forward almost two decades before the debate on perceptual content was clearly defined, Armstrong's belief theory of perception, developed in two of his early books, Perception and the Physical World (1961) and A Materialist Theory of Mind (1968) can nonetheless be shown – in light of Gareth Evans' introduction of "nonconceptual content" – to straightforwardly imply a "conceptualist account" of perceptual content in the sense of that phrase found in this literature.

thing, *a*, with the content *a is F* and *a is not-G*.[46]

This principle implies two things about conceptual content. First, that it is consistent. So for instance, one could not hold two contradictory beliefs concerning a that it is both F and *not*-F. If this were possible, it would follow from the Principle of Cognitive Significance that F is different from itself, which is clearly absurd. The principle also highlights the *subjective* character of conceptual content under the Fregean framework. Recall that the *sense* of a content was partially fixed with reference to the cognitive states of the subject – beliefs, judgements, etc – which suggests that the subject's "epistemic capacities" effect what his experiences and thoughts are properly about (Hyde, but not Jekyll, for instance).[47]

C. The Principle of Reference Determinacy runs as follows:

If content c is conceptual, then a subject, *s*, can determine the semantic value of c.[48]

This can mean a number of things; knowing semantic value can be a matter of having the ability to classify, identify, recognize and re-identify the referent of *c*. If one were to fail to have one of these abilities regarding a certain state it would

---

[46] Ibid. p. 11

[47] Ibid. p. 11

[48] Ibid. p. 12

be hard to see why we should attribute conceptuality to the content of the state being considered. So, for instance, an inability to identify any bird *as a* bird beyond *this bird* here in front of me would mean, in effect, that I have no concept of bird whatsoever. Alternately, if I could not, at a later time, re-identify a referent as being an instance of a concept, then I do not possess that concept either. For example, I might claim to know what a "shaftail Finch" is as I observe it on its own, and yet fail to pick it out of a line-up of similar looking species of Finch at a later time. It seems that my failure can only be due to my not really possessing the concept "shaftail Finch" in the first place. Re-identification in particular will loom largely in our discussion of demonstrative concepts and unnamed colours.

Reference determinacy also helps to make sense of how we could be in error concerning the way in which we represent the world. Since the determination of a reference is undertaken by a finite individual with imperfect knowledge of the external world, our representation of the world – a representation that might seem to us perfectly reasonably – can often fall short of reality or otherwise go awry. Recall that the ability to represent incorrectly is a necessary component of representationalism; a state which only exhibited a law-like connection to the external world is not a perceptual state but merely a sensation.

The principles Gunther offers cover several different features that are commonly associated with conceptual content: expressibility, communicability,

productivity, subjectivity, the capacity for misrepresentation, and so on. While certainly not exhaustive, these three principles will see us through the discussion.

One final note: It might seem from all of this that whether or not all experience is conceptual or nonconceptual is a highly specialized, localized and otiose debate lacking any real consequence for all but a handful of researchers. It will be worth taking a moment to dispel this thought. As Christopher Peacocke points out, the question of conceptual content is inseparable from many important issues. Working outward from those issues which would concern mainly theorists to those which would interest most philosophers, the conceptual/nonconceptual content debate has important things to tell us about: the individuation of conceptual content, the nature of concept possession, the nature of rationality, our conception of objectivity, the relation between animal and human perception, and how widespread "feelings" are in the universe.[49] Consider this last issue, that of "feelings." Suffice it to say that non-conceptual accounts will tend to characterize phenomenal perception as a construction of perceptual capacities culled from phylogenetically more ancient (one might say, fundamental) modes of perceiving, with our conceptual abilities lain overtop of perception. There is a certain continuity in the conscious perceptions from higher animals and pre-linguistic

[49] Peacocke (2001) p. 239

infants through to language using adults: they are all phenomenally conscious, and so they all have feelings. But whether conceptualists will extend phenomenal perception to infants and animals will depend on how stringent their requirements for concept possession are: whereas on John McDowell's account, which presumes that concept use is only possible for linguistic creatures, will deny this claim others like Alva Noë, who believes that concepts extend beyond language use, will affirm it. But consider the implications of McDowell's position. It would follow that pre-linguistic infants, great apes, and everything else on down the Great Chain of Being are not phenomenally conscious. There would be no "light" on inside any of these creatures. Proof (or disproof) of such a thesis would obviously have important implications for work in fields like animal ethics, not to mention that, *existentially*, it would be on par with, or even worse than, discovering that there was no life anywhere else in the universe.

2. Conceptualism

"This conception and belief which nature produces by means of the senses, we call perception."

- Thomas Reid[50]

2.1 Elucidation

David Hamlyn presents a useful analysis of conceptualism in terms of the distinction between perception and the causal processes which underlie it. The earliest theories of perception, he says, seemed to have been purely causal; an assimilation of perception to sensations in which perception was entirely the result of an interaction between the sense-organs and their objects. Hamlyn relates Aristotle's theory of vision as an example of this sort of causal theory. For Aristotle, the organism is made *sensitive* to its surroundings when its sense-organs receive or extract the *form* of objects without receiving the *matter* of which they are comprised.[51] Once the sense organs have taken on the form of the things in the world, they must then "transmit" this form to a faculty Aristotle called the *sensus communis,* a sort of "master organ" responsible for binding the various sensations

---

[50] Reid, *Essays on the Intellectual Powers of Man* p. 210

[51] Aristotle, *De Anima* (424a25-9)

of the world together into a coherent *perception*.[52] This coherent perception is, so to speak, that which is placed before the mind.

Aristotle understood perception as a passive, physiological achievement in which the influences of our active faculties of judgement and imagination were entirely absent. He seems to rule out the possibility that perception might involve judgement on the grounds that judgement, unlike perception, has no canonical contents. Whereas the faculty of perception, the *sensus communis*, can have nothing as its object other than those things that are sensible (those qualities heard, touched, smelled, etc.), the faculty of judgement is not bound to consider any one content. The faculty of judgement can consider perceptions, but it can also consider imaginings, numbers, the sensations which make up perceptions, and even itself. The *sensus communis* must therefore be, in some important sense, "unmixed." The immediate object before our mind is only comprised of the impingements of the world itself on the sense-organs. We might be misled by the mention of the sensus communis into thinking that this might be something like the Kantian understanding, but it seems to be something more like a path onto which the sense-impressions from each modality merge rather than an apparatus which fundamentally changes the incoming sensations.

---

[52] Bennet and Hacker (2003), p. 17

Hamlyn goes on to say that such views have been demolished in light of the development of representational theories of perception, which made a distinction between sensations, which could be explained in terms of physiology, and perception, for which we needed a "mental" level of explanation. Of course, representationalism does not deny that there are in fact sensations which arise out of a causal connection between our bodies and objects, but rather only asserts that the immediate objects of our perceptions are mental representations. Now we might think, having gotten as far as representationalism, that we would have be faced with a choice between the three options I outlined at the beginning of the last chapter: conceptualism, nonconceptualism, or a mixed content theory. But conceptualists feel that the affinities between states with "nonconceptual content" and "sensations" make the latter two explanations impossible. Even if we were to find a way of making sense of a representational state bearing "nonconceptual content," we would be mistaken if we tried to apply it to instances of phenomenality. At best, it might help us to elucidate certain sub-personal exchanges of information – say, between brain and world – but this will not bear on how the organism immediately perceived the environment.

Conceptualists feel that appeals to sensations and nonconceptual content are inadequate for explaining the content of perception since perceiving involves a level of interpretation of the perceptual field which necessarily employs the use of

concepts. Normally, or *naturally*, when I look at the world I see various objects, and not merely an array of undifferentiated sensations. That perception is dependent on concepts is indicated, as Hamlyn claims, by the fact that "in order to perceive something we must perceive it *as such and such.*"[53] All perceiving is perceiving-as.

With this in mind, let's turn to an analysis of an actual conceptualist theory. I will take John McDowell's arguments for conceptual content as the focal point for a discussion of conceptualism, since he is the most well known and widely read of the conceptualists. His analysis is therefore the preferred starting point for both a defence of and an attack upon that position. McDowell's influence on conceptualists such as Bill Brewer and Alva Noë is obvious, and their analyses can usefully be described as elaborations on McDowell's initial position. On the other hand, nonconceptualists such as Christopher Peacocke, Jose Luis Bermudez, and Sean Kelly, each of whom will be part of the discussion on nonconceptualism, engage McDowell's theories and offer rebuttals to McDowell's criticisms of nonconceptualism as well as further arguments against conceptualism.

In his influential book Mind and World, McDowell has developed a

---

[53] Hamlyn (1994) p. 91

conceptualist account of perceptual content. In this chapter, I will summarize John McDowell's position. I will then turn to some of the many critiques of his theory, as well as to McDowell's rebuttals to these critiques. I will conclude that McDowell's theory of perception is incomplete. Like virtually all conceptualists, McDowell's conception of the role of perception is simply too narrow; he focuses solely on its epistemological function while neglecting its role in activities such as the shaping of motor action. This incomplete picture of perception in turn informs his position on perceptual content, so that McDowell's appeal to conceptual content helps to make sense of certain perceptual abilities, such as those associated with reasoning and object recognition, while falling short of explaining certain other features of our perception, such as its fine-grained nature (despite McDowell's attempt to meet this challenge) as well as other of our perceptual abilities, such as our capacity for visually-guided action.

Like many conceptualists, McDowell's motivations for adopting conceptualism are epistemological. He is interested in developing a picture of how our judgements and beliefs can be grounded in sense experience or, put differently, how experiences provide us with reasons for our beliefs. As we will see, for McDowell a coherent picture of how experience can stand in this relation to beliefs can only be accomplished if we abandon certain dogmatically-held distinctions between cognition and perception, or between thinking and sense,

which have the effect of making this relation impossible. McDowell focuses his critique on the dual prejudice involved in conceiving cognition as "spontaneous," involving routines which have an element of freedom ("of thought," we might say), while regarding perception as "passive," where we are only the recipients of perceptual states and not their authors. In place of this dichotomy he urges that we instead understand perception itself as necessarily employing the capacities for spontaneity which are most naturally associated with thought. As we will see, for McDowell nonconceptualism makes an adequate explanation of our epistemological practices impossible, since it maintains this absolute distinction between our "spontaneous" mental states, such as judging and believing, and our "receptive" perceptual states, such as hearing and seeing while insisting that it is the *latter* which makes up the look of experience.

A. Logical spaces

McDowell borrows terminology from Wilfrid Sellars in order to make his line of reasoning explicit: In Sellars' analysis, perception and reasoning exist in different "logical spaces," respectively, the *space of nature* and the *space of reason.*[54] If perception is understood in this way then it is difficult to see how it can be said to supply us with "reasons" which could justify our beliefs and

[54] McDowell (1994) p. xiv-xv.

judgements. Let's briefly consider both of these terms and their relations in turn, since they are central to McDowell's theory.

i. The logical space of reasons consists of "things" (to be non-committal for the moment) which can stand in relations of warrant or justification to each other. Among such "things" are, for instance, my belief that I am overdue for a flu shot, which is justified by my further belief that it has been two years since my last shot, which in turn is justified by certain memories I possess of the relevant events, while my reliance on these memories is itself warranted by my (likely unconscious) belief that the memories I have are of events that actually occurred, and so on. The set of "things" in the space of reasons we will be concerned with here are cognitive states, broadly conceived: beliefs, judgements, etc. The essential attribute which all cognitive states share as members of the space of reasons is that they justify or are themselves justified by further cognitive states, so that, for instance, we can give *reasons* for believing certain things by referencing further beliefs. In McDowell's words: "The topography of the conceptual sphere is constituted by rational relations."[55]

Another characteristic of the space of reasons is its *spontaneity*. This simply means that thoughts are not bound by natural laws to unfold in some fixed

[55] Ibid. p. 5

manner. Relations of thought are not impressed upon us in a way that is beyond our control. Thoughts seem – at least from the perspective of the thinker – to exhibit some autonomy from *the way things are* (we might also say they are stimulus-independent). In other words, the faculty of thinking (what Kant called the understanding) is spontaneous in that it "is exercised in actively self-critical control of what one thinks, in light of the deliverances of experience."[56] McDowell also equates spontaneity (and therefore, reason) with the notion of freedom, and speaks of conceptual capacities as belonging to a "faculty which empowers us to take charge of our lives."[57] The space of reasons is the realm of freedom. Note the similarity here between McDowell's characterization of the space of reasons and the Generality constraint for concept possession outlined above. That my thoughts enjoy autonomy from the way things are means that I have the freedom to employ any concept in any thought whatsoever.

ii. The logical space of nature, on the other hand, is "the logical space in which the natural sciences function"; it is the realm of law. "Things" within the logical space of nature are connected to one another through causal (or at least law-like) relations. To take a well known and non-mental example: the number

---

[56] Ibid. p. 49

[57] Ibid. p. 73

of rings appearing on the trunk of a tree is causally covariant with the number of years the tree has been growing, so that, presuming that normal environmental, phylogenetic and ontogenetic conditions hold, a seven year old tree will exhibit seven rings.

McDowell does not offer much more in the way of a positive description than this, since what is most important about the space of nature for our considerations is that it does not share the other features of the space of reason: "The relations that constitute the logical space of nature, on the relevant conception, do not include relations such as one thing's being warranted or – for the general case – correct, in light of another."[58] For instance, it surely wouldn't make sense to say that the Moon's gravity *justifies* or *warrants* the ocean's tides, but simply that it *causes* them. This is, and can only be, an empirical description; it lacks the *normative* dimension which Sellars attributes to the relations of *justification* between thoughts.

These may appear to be uncontroversial, almost trivial, characterizations of the relations involved, respectively, in "mental" events and in "worldly" events. And yet McDowell suggests that when we examine the interface *between* the mental and the worldly we run up against a problem that has been the cause of

[58] Ibid. p. xv

much "philosophical anxiety" and which gives rise to questions like "How is it possible for there to be empirical *knowledge*?" or, more generally, "How is it possible for there to be thinking directed at how things are?"[59]

B. Antinomy

It is a standard claim of what McDowell calls "minimal empiricism" that our beliefs can be verified or disproved by experiences of the world impressing itself upon perceiving subjects.[60] But how can this be, if judgement and experience exist in two different logical spaces, each with its own sorts of relations? It seems as if our judgements and our experiences are separated by an unbridgeable gulf which bars us from having any sort of genuine knowledge about the empirical world. This anxiety forcefully arises when we consider the freedom which the play of concepts has within the space of reasons. Spontaneity, which is a key characteristic of the space of reasons, is an idea of freedom which threatens to "make what was to be empirical thinking degenerate" so that we are left with what McDowell evocatively calls "a frictionless spinning in the void."[61] If we enjoy a complete freedom of thought so that nothing from outside the conceptual

---

[59] Ibid. p. xiii

[60] McDowell (1994) p. xvi

[61] Ibid. p. 66

sphere can constrain the play of concepts, then the very possibility that judgements which are about the empirical world are grounded in experiences of the real world is threatened.[62] So the line of reasoning which began with the sensible division between "reasons" and "nature" comes up against a sort of Kantian antinomy in which it is obvious to us that the deliverances of experience *must* somehow act as a rational restraint on our judgements and yet at the same time it seems "hard to see how experience could function as a tribunal, delivering verdicts on our thinking."[63]

This analysis forces us into a position in which our conceptual scheme is set at a distance from our experience, so that any attempt to ground our judgements seems to be an impossible task. Exercising concepts would therefore seem never to amount to empirical knowledge, since judgements would have no necessary relations to experience. In order to assuage the philosophical anxieties this problem raises we must either bite the bullet and insist that experience must rationally restrain our beliefs even though it is inconceivable to us how it might play this role, or instead retreat into the view that beliefs can only ever be justified by further beliefs and not by an "impact" from the external world.

---

[62] Ibid. p. 5

[63] Ibid. p. xii

C. Dilemma

McDowell considers two possible solutions that would assuage this anxiety, one which would be happy to leave experience out of the "reason-giving" picture altogether so that the coherence of the beliefs themselves are their only rational constraint, and the other which wishes to extend the space of reasons so as to accommodate brute sensory impacts from the outside world, which would then serve as a form of rational constraint on judgements. The former, endorsed by Donald Davidson, is a form of coherentism, while the latter is an appeal to "the Given." But for McDowell neither option is desirable; coherentism and "the Given" for him form two horns of the dilemma which arises from our initial division of logical spaces. Neither position gives a satisfactory account of the role of experience in knowledge; while coherentism simply gives up on fitting experience in as a rational restraint altogether, those who appeal to "the Given" assert that experience can constrain our beliefs rationally without offering a satisfactory account of how this is possible. I will now examine both of these positions and McDowell's criticisms of them a little more closely.

i. The Given – The most obvious reply to the worry that our beliefs might not have a rational restraint is to suggest that *experience* can fulfill this role. All of those beliefs and judgements which are in some way about the empirical

world must be answerable to experience. So on this view, although experience is brought about by law-like connections with the outside world, an attribute which would seem to place it in the space of nature, it nonetheless exists *inside* the space of reasons since it can be called on to justify certain of my beliefs. In fact, it is the law-like connections which make up perception that render it useful not only in justifying but also in *rationally restraining* my beliefs. It is only because perception is governed by law-like connections that it can serve as a stable ground on which to base my epistemic practices. Perceptions simply *come at me* and I am unable to choose or change willfully what I will see when I open my eyes.[64] And, on this picture, in justifying an empirical belief I can simply *point* to a bit of the conceptually unstructured experience with which I am saddled.

This position which I have briefly outlined makes an appeal to what is often called "the Given." As with the division of logical spaces, McDowell borrows the concept of the Given from Wilfrid Sellars, whose canonical paper "Empiricism and the philosophy of mind" contained a sustained criticism of that concept. The idea of the Given, says McDowell, "is the idea that the space of

[64] Of course, we do choose what we see in some respects. We choose, for instance, where to look, we can control the amount of light that enters our eyes by squinting, etc. McDowell recognizes this, and responds: "But one's control over what happens in experience has limits: one can decide where to place oneself, at what pitch to tune one's attention, and so forth, but it is not up to one what, having done all of that, one will experience." [McDowell (1996) p. 10]

reasons, the space of justifications or warrants, extends more widely than the conceptual" and that this extra space "is supposed to allow it to incorporate non-conceptual impacts from outside the realm of thought."[65]

The purview of rational relations is extended to that which is given in experience, which can then act as the substratum upon which all knowledge is founded. The Given is therefore a critical concept in most versions of the epistemic project of *foundationalism,* which is the attempt to analyse knowledge as arising out of a connection between our epistemic practices and some ultimate justificatory ground. To speak loosely, if we can *reach outside* our conceptual scheme and take hold of the world *as it is*, occurring out there on the outskirts of the space of reason, we will have some way of fixing our knowledge to something concrete. If, on the other hand, we only had access to our own conceptual schemes in our search for knowledge, if our conceptual scheme did not have 'experience' upon which to gain *traction*, then our claim to empirical knowledge, and further even our claim to empirical thought *in toto* (of which empirical knowledge is an instance) would be jeopardized. So it seems to follow from there being a Given in experience that judgements which are "hooked up" to the Given in just the right way would constitute empirical knowledge, since perception is

---

[65] McDowell (1996) p. 7

assembled from law-like processes and, in that sense, can put a reign on the spontaneity of thinking..

The results of this picture for observational concepts are especially important. An observational concept – by which I mean some concept that explicitly involves a perceptual identification: a demonstrative concept such as "that shade" is an example – grounded in the world as it is in virtue of its connection to the Given. And from there, once we have captured a piece of the Given in some conceptual form, it would seem to be a straightforward story, since there is no trouble in deploying the observational concept in service of rationality.

However, both Sellars and McDowell argue that the Given simply cannot do the work foundationalism saddles it with. Events which take place in the space of nature, even those events which are impressions on our sense organs, *cannot* stand in a relation of justification or warrant to another state. It follows that perceptually-based beliefs – say, about the objects in your immediate environment – cannot be justified by the experiences which they are putatively about. Supposing that we *could* take experiences – construed solely as impacts on the brain from the outside world – to justify beliefs would be a case of taking it that "empirical description can amount to placing things in the logical space of

reasons."[66] An observational concept cannot be part of a rational system since it remains only a description of a thing in the space of nature.

But this might seem to be a mere contradiction of the idea of Givenness and not a refutation. So here is another, clearer way to put it. Recall that the Myth of the Given presupposes the neat division of logical spaces which was outlined above, in which sensations, since they reside in the space of nature, are events which we can only *passively* register and are therefore beyond my ability to alter in spontaneous and self-critical ways. The aspect of the Given which McDowell finds objectionable is "the idea that what we perceive or experience is the *passive* end product of the stimulation of a specific faculty, namely the physiological apparatus of the senses."[67] No other faculty, especially one of a conceptual or inferential nature like Kant's "understanding," contributes to the perceptual process, and so the Given is not therefore conceptually or inferentially structured in any way. But since we only passively register the Given – since it is a sort of "brute impact" from an exterior where the free play of our concepts has no dealings and therefore has no part in its formation – it is incoherent to claim that we can nonetheless appeal to it in explaining how our judgements are

---

[66] Ibid. p. xv

[67] Roy (2003) p. 18

justified. At most, we might be able to say that we are not to be *blamed* for the beliefs which we come to have based on our perceptual experiences; they are forced upon us, after all! But that we are blameless since we are compelled by a brute force into a certain position is far from constituting a justification for our position. As McDowell puts it in several places, the Given only "offers [us] exculpations where we wanted justifications."[68]

ii. Coherentism – The attempt to extend relations of justification to the impacts of a Given with which we are passively presented has failed, and its failing was a result of the illicit move in trying to smuggle perception across the border between passivity and spontaneity, cause and reason. Donald Davidson regards this result as a death knell for empiricism. Adding to Quine's store, he refers to the Myth of the Given as a "third dogma" of empiricism, and, it stands to reason, one that it simply cannot do without; even a "minimal empiricism" becomes untenable if we deny the role of experience as a rational restraint on empirical beliefs. Davidson's position therefore begins with the claim that although experience is obviously *causally* relevant to one's beliefs and

[68] Ibid. p. 7. In a footnote, McDowell gives this helpful illustration: "[If] someone is in a place from which she has been banished, she is *exculpated* by the fact that she was deposited there by a tornado. Her arriving there is completely removed from the domain of what she is responsible for; it is not that she is still responsible, but there is a basis for mitigating any sanctions."

judgements, it has no part in conferring *justification* to those mental states. Experience therefore cannot provide us with *reasons* to believe anything whatsoever.

Having accepted the results of Sellars' attack on the Given, Davidson retreats into a form of coherentism where "nothing can count as a reason for holding a belief except another belief."[69] The Given is left out of the picture altogether so that the only justification we can have for a belief will be the way that it hangs together with all of our other beliefs. In a sense, the result is that our concepts *are* in fact a frictionless spinning in the void, having no external restraint; it's just that the demand for an external restrain was misguided to begin with, since it could only be filled by something which would illicitly straddle the rational/causal divide.

But coherentism, says McDowell, gives rise to philosophical anxiety, since it confines us within the realm of thought so that we are not *in touch* with anything outside of it; our bodies may be causally affected, but our minds remain, in a sense, untouched by the external world. But this makes it hard to see how coherentism can actually account for the beliefs which we have of the empirical world, or whether it instead leaves us with no convincing way to credit ourselves

---

[69] Davidson (1986) p. 310

with empirical knowledge at all; in other words, with thoughts emptied of content. I will not be concerned here with whether or not there *is* some way to satisfactorily meet this challenge within a Davidsonian framework, but I will simply note McDowell's skepticism when he suggests that Davidson does nothing to allay the worry except to blandly suggest that we can in fact get along with merely having consistent beliefs: "Davidson manages to be comfortable with his coherentism, which dispenses with the rational constraint on thinking from outside it, only because he does not see that emptiness is the threat."[70]

So the antinomy remains. It is not solved by either approach, since we end up merely oscillating between these two possibilities: On the one hand, we recognize the absurdity of believing that a mere causal connection to the world could serve as justification for our beliefs, and so an appeal to the Given as a sort of bare presence is hopeless. So we are left, by default, with coherentism. On the other, the consequences of coherentism seem to make empirical thinking impossible altogether since it puts thought at a distance from reality, an alarming result which naturally draws us back into an acceptance of the Given, despite its difficulties. Our return to the Given is spurned on by our hope that "the spontaneity of the understanding can be constrained by the receptivity of

[70] McDowell (1996) p. 68

sensibility," so that we can thereby be "entitled to the very idea of subjective postures with objective import.[71]

But, says McDowell, the solutions proposed by both Davidson and the proponents of some form of Givenness are blinded by their dogmatic separation of the attributes of the space of reasons from the space of nature to another way of rescuing a minimal empiricism and thereby easing our philosophical anxieties, a way which McDowell himself put forward.

D. The Neo-Given

I call McDowell's answer to this philosophical anxiety which coherentism and the Given have failed to assuage a "neo-Given" theory, since he offers no title for it himself. However, it is not an entirely arbitrary choice on my part; McDowell's solution depends on augmenting our understanding of the Given in order to show how it *can*, despite the failure of its classic rendering, act as a rational restraint on our beliefs. This results in a picture of experience – a synonym for the Given, we might as well say – which is different in significant ways from the one which was debunked by Sellars and branded a hopeless theory by Davidson.

In a way, we can understand McDowell's project as an attempt at a

---

[71] McDowell (1998) p. 366

thoroughgoing naturalism, in that it tries to put back together what the separation of logical spaces pulled apart. The divide between logical spaces resulted in a separation of duties: science delivered its understanding in a logical space in contrast with the space of reasons which is (perhaps) the workspace of the epistemologist and his ken. Now, it is tempting in light of this separation to think that "the distinctive features of [the logical space of nature] capture the very idea of what is natural" so that what is natural becomes synonymous with what is non-mental, that is, among other things: publicly observable, quantifiable, exhibiting law-like relations, etc.[72] But though there is a sense in which our conceptual capacities *do* contrast with the workspace of modern science and its distinctive kind of understanding, it does not follow that concepts, relations of warrant and justification, etc. are somehow *unnatural.*

McDowell urges us instead to conceive of mental states which involve conceptual capacities, such as beliefs and judgements, as part of the "second nature" of their possessors. This means, on the one hand, that we should not conceive of the mind as somehow separate from nature: a standard enough thesis. But it also implies that a "thing" which we take to be in the space of nature – and he obviously has experience in mind here – *can be conceptually structured,* since

[72] Ibid. p . 367

concepts are no longer conceived as something outside the purview of the natural.

That is McDowell's solution to the antinomy above, and also happens to be the central thesis of his book. He wishes to preserve a minimal empiricism – and therefore the role of experience as a "tribunal" – but is sensitive to the difficulties inherent in the Given. He recognizes the merit in Davidson's point that any impingements across the outer boundary of concepts could only be causal and not rational. So he simply moves the border of the conceptual *out*, to encompass the outer edge of experience. That way experience is no longer something which must be smuggled across the conceptual divide into rational relations, since *experience itself* is now conceptualized. McDowell puts the matter this way:

> **It is not that [conceptual capacities] are exercised *on* an extra-conceptual deliverance of receptivity. We should understand what Kant calls "intuition" – experiential intake – not as a bare getting of an extra-conceptual Given, but as a kind of occurrence or state that already has conceptual content.**[73]

Now that experience is understood to bear conceptualized content it can be rationally linked with even concepts which are very remote from any empirical observation, through the system with which we continually adjust our judgements to the deliverances of experience. But it is important to realize that McDowell is

---

[73] McDowell (1996) p. 9

not suggesting experiences and judgements are therefore one and the same, both being free and spontaneous. If McDowell's system were to collapse the distinction between experience and judgements, so that neither could be said to be passive, then he would be leaving us no better off than did Davidson: with no ground on which our concepts can gain friction; with a coherentism lacking any external restraint. This would be absurd, since it would mean that we could have some kind of spontaneous control over the content of our perception, which we clearly lack.

With this new rendering of the Given, says McDowell, we can finally make sense of how empirical knowledge is possible. His answer is Kantian, though he differs importantly with *at least certain readings* of Kant on an important issue. Kant believed that thought arose out of a *cooperation* between two faculties. The faculty of sensibility gives us "the capacity to acquire representations through the way in which we are affected by objects," that is, it provides us with a content or basis for thought.[74] It operates in a way that is wholly passive, or receptive, and so the character of the representations it produces do not depend on the operations of understanding, which exhibit spontaneity. But for a mental state to become a *thought* requires, in addition to

[74] (A19/B33)

sensibility, the faculty of the understanding, which is the "the logical, discursive, and proposition-forming capacity of the mind, which produces concepts as outputs".[75]

I stress that McDowell's position is different than *certain* readings of Kant because we can also read Kant in a way which is much closer to the spirit of McDowell's analysis. On some readings of Kant, the above analysis is taken to mean that *experiences* necessarily employ concepts, so that a bare intuition would be *nothing at all* for the creature subject to that intuition, but only a "sub-personal" accomplishment of the *faculty* of intuition. As eminent a Kant scholar as P.F. Strawson seems to endorse this position:

. . . we can form no conception of experience, of empirical knowledge, which does not allow of our becoming aware in experience of particular items which we are able to recognize or classify as instances of general kinds or characteristics.[76]

This sounds very much like McDowell, who, in a discussion of animal perception, says:

In the account of inner workings, one sub-froggy part of a frog transmits information to another: the frog's eye talks to the frog's brain. In the sense in which the frog's eye tells the frog's brain things, nothing tells the frog anything.[77]

---

[75] Hanna (2005), p. 192

[76] Strawson (1966) p. 47-48

[77] McDowell (1994) p. 197

If we understand "the frog's eye telling the frog's brain things" as a matter of sensing, or of taking in a bundle of information from the external world, then McDowell is here claiming that the content of a sensation is different in kind, and in its intended target, from the content of a full-blooded thought. The receiver of sensational content is not the animal itself, but only some part of the animal's physiology, and we run the risk of massive error – viz. the Myth of the Given – if we conflate these two stories. The sub-personal story about the sensory system, in terms of an information-processing device communicating to other parts within the animal, has nothing to do with the story about what the sensory systems are like *for the animal itself.* Therefore, receptivity, understood as a brute, *nonconceptual* impact from the external world on the animal's senses, does not make an even notionally separable contribution to the co-operation between the understanding and sensibility which gives rise to experience.

2.2 *Critique*

I grant McDowell one of his major theses: only states with conceptual content can stand in rational relations to our beliefs and judgements. In fact, I am endorsing a mixed content view of phenomenal states in part because I think that we can only make sense of experience justifying beliefs if we allow that some experiences are conceptualized. In any case, the other two possibilities don't

seem to be very promising.

On the one hand, we could bite the bullet and accept that experiences *don't* justify beliefs. As I have already said, I do not have the room here to fully evaluate coherentism, except to say that, like McDowell, I find the position unfavourable. Nothing seems clearer to me than that my belief that "There is a dog on a mat in front of me" is justified by my experience of the dog on a mat.

On the other, we could try to tell a story about how a nonconceptual experience *can* justify beliefs. But since the Given isn't conceptually structured this seems hopeless. So I endorse *most* of what McDowell has said in his critique of the Given and nonconceptual content. Appeals to the Given cannot justify our beliefs. In fact, I am sympathetic with C.I. Lewis' claim that we cannot directly appeal to the Given at all, and that what we can articulate of our experiences is always already conceptual:

> **There are not concepts of immediate qualia as such [. . .] because articulation is the setting of bounds and establishing of connections; because what does not affect discrimination and relation has no handle by which the mind can take hold of it.**[78]

But what about statements such as "That ineffable shade of yellow is the same shade that I saw last week"? This seems to be an example in which a nonconceptual feature of experience – "ineffable shade of yellow" is used in

---

[78] Lewis (1929) p. 128.

making a judgement. But while the phrase "ineffable shade of yellow" might *approximate* my Given experience of a shade, insofar as I have labelled it I have already brought it under a concept which does not fully capture the experience. In exchange for the ability to deliberate upon the colour, to make judgements about it, I have sacrificed the informational richness of the initial perception. For instance, I gloss over those properties of the colour (whatever they may be) which were discriminated by my visual system but which I did not notice myself. The extent of the informational loss is something like the difference between describing the sight of a coffee cup sitting on a table to you and showing you a picture of it. To cite Lewis once more:

> **[There] is no knowledge merely by acquaintance; . . . knowledge always transcends the immediately given. The merely contemplated or enjoyed may possess esthetic significance, but if it is to have cognitive meaning this immediacy must become the subject of an interpretation which transcends it[.]**[79]

So my issue is not with McDowell's claim that only conceptual content can enter the space of (at least *rational*) reasons, but rather with his claim that it follows that *all we need* is conceptual content. I will attack this aspect of McDowell's theory on three fronts. First, I will suggest that nonconceptual content can be useful in explaining how we come to possess concepts in the first

---

[79] Lewis (1929) p. 118-119

place. Next I will give several examples of conscious perception that seem to be concept-independent. Drawing on the work of Sean Kelly I will make the case that conceptualism fails to account for certain features of perception and so it cannot be a complete theory of phenomenal experience. We can have perceptual states which are not necessarily specified by conceptual content, and this is a strong motivator for positing a kind of nonconceptual content. Lastly, I will criticize two claims McDowell makes about the explanatory value of conceptual content: that the self-world distinction which is a necessary prerequisite for experience can only be achieved by conceptualized content; and that the world itself can only be parsed into things, places, etc. with the use of concepts. I think these positions are wrong; both of these roles can be filled by nonconceptual content, which I hope to show by appealing to arguments made by Jose Luis Bermudez, Andy Clark, Adrian Cussins, and Susan Hurley.

Since McDowell's system is built around the assumption that we must conceive of perception as conceptualised for it to count as experience and for it to have rational relations with our epistemological practices, an exposition of how these conditions can be met by nonconceptual states would undercut the primary motivation for conceptualism. Of course we can still make a case for there being a conceptual content in experience. But although invoking conceptual content may help in accounting for features of both our phenomenology and our

epistemological practices, I believe it follows from these following arguments that it cannot be the whole story.

A. Concept-independent perception

i. Concept genesis – A criticism that is levelled against nonconceptual content is that it is hard to see how it helps us to parse the environment into distinct objects, planes, properties etc. Peter Carruthers remarks that nonconceptual content might be chimerical, since "there can surely be no *representation* without *discrimination* . . . yet all discrimination might be thought to implicate concepts."[80] As we will see in the discussion of *affordances* below, there are good reasons for rejecting this claim; animals perceive distinct characteristics in their environment which afford them certain opportunities for action without at the same time having to employ concepts to make these discriminations. But I believe there is an even more powerful consideration for believing that discriminations outstrip conceptual abilities: if they didn't, if all discriminations were accomplished conceptually, then how do we initially master these concepts? As McDowell would have it, all our dealings with the world are mediated by concepts so that even the most rudimentary experience directly enters into the 'commerce' of our rationality. But however attractive the neo-Given

[80] Carruthers (2000) p. 130

theory might be in solving epistemological problems, it falters in giving an account of how these conceptual abilities could first arise in experience.

The only explanation for this which is open to conceptualism will be viciously circular: conceptual content arises from conceptual content. For instance, you might use the concept "jagged" to approximate the shape of a mountain range off in the distance, and this concept will be grounded in the experience which is itself fully specified by the demonstrative "that shape." But both of these concepts are equally context-free, recombinable, general, etc; so there will be no explanatory value in citing the more general concept "that shape" in an explanation of how we come to master the concept "jagged." Or imagine trying to give an account of how one came to master the concept "that shape." It seems that a conceptualist could only appeal to the experiences in which "that shape" is applicable. But since this experience already *contains* the concept "that shape" this amounts to no explanation at all; as Peacocke remarks: "An account of mastery of a concept is still circular if it adverts to the enjoyment of perceptual states with a content requiring possession of the concept whose possession was to be elucidated."[81] So *if* we think it is important that we should have a genesis story about how we can come to master concepts, then we will have reason to doubt that all experience is

---

[81] Peacocke (1992) p. 115

conceptualized.

In order to answer this challenge it is reasonable to assume that there *must* be some manner of discrimination which doesn't involve concepts in order for us to be able to account for how we come to master concepts in the first place. One of the obvious advantages of the notion of nonconceptual content is the promise it offers for explaining what it is to possess concepts and how to acquire them. Much like primitive forms of self-awareness can act as a springboard to full fledged ascriptions of "I" thoughts, which we will be discussing shortly, nonconceptual contents of experience can be the opportunity for devising general concepts which are based on the features of that perception.

In order to see how this works Peacocke asks us to consider the case of mastering the concept "square." He suggests that we can get around complications in the conceptualist account of concept mastery if we accept that the nonconceptual *way* in which something is experienced makes available various different demonstrative concepts which can be used by a thinker to specify – albeit imperfectly, as we have seen – the content of that experience. Peacocke suggests that there are several ways of capturing the "way" a shape is perceived using different demonstratives. So for instance, every time I perceive a shape with four sides of equal length with one of its sides laying upon a plane, I will feel "primitively compelled," as he puts it, to apply the term "square" to that

experience. But it should not be assumed that my perceiving that shape in that "way" was dependent upon my possessing the general concept "square." The way in which I perceive the shape does not involve concepts, but only properties and relations such as "curved, parallel to, equidistant from, same shape as".[82] For instance, my experience of a square as opposed to say, a diamond shape is fixed by my perceiving one or the other sets of symmetries:

> **When something is perceived as a diamond, the perceived symmetry is a symmetry about the bisector of its angles. When something is perceived as a square, the perceived symmetry is a symmetry about the bisector of its sides.**[83]

Experiencing one or the other does not depend on having concepts such as "square" or "diamond."

Another example of a nonconceptual origin for a concept comes from Charles Taylor. An oft-cited passage reads:

> **Our perceptual field has an orientational structure, a foreground and a background, an up and down. . . . This orientational structure marks our field as essentially that of an embodied agent. It's not just that the field's perspective centres on where I am bodily – this by itself doesn't show that I am essentially agent. But take the up-down directionality of the field. What is it based on? Up and down are not simply related to my body – up is not just where my head is and down where my feet are. For I can be lying down, or bending over, or upside down; and in all these cases up in my field is not the direction of my head. Nor are up and down defined by**

---

[82] Ibid. p. 119

[83] Ibid. p. 118

certain paradigm objects in the field, such as the earth or sky: the earth can slope for instance. . . . Rather, up and down are related to how one would move and act in the field.[84]

The meaning of our concept "up" appears to be situational in a way that is not shared by concepts like "bachelor," for instance. I believe that both of these examples establish that the possession of certain of our concepts – presumably, those that are most closely tied to direct observation (such as "diamond-shaped") and action (such as "up") – is best explained, as Adrian Cussins says, "in terms of certain basic, nonconceptual abilities that we possess, such as our ability to move and act in a co-ordinated way."[85]

ii. Situation dependence – Sean Kelly's criticism of conceptualism centers on the claim that properties are not, as presented in experience, independent of the contexts in which they are perceived or the objects they are perceived to be properties of.[86] In the case of colour for instance, Kelly is suggesting that the experience of colour cannot be abstracted from the scene they inhabit or the characteristics of the object in which they inhere. In light of this fact about colour experience, Kelly argues that demonstrative concepts such as "that colour" cannot specify these sorts of dependencies, and are therefore

---

[84] Taylor (1978-1979) p. 154

[85] Cussins (1990) p. 146

[86] Kelly (2001) p. 227

inadequate for fully specifying the content of a perceptual experience.

Kelly uses the phenomena of colour constancy to exhibit a colour's dependency on its *context*. A surface is said to exhibit colour constancy if it retains its apparent colour under different conditions of illumination which alter the way the colour is experienced. For instance, when I look at the door to my room it seems to be a uniform shade of white, although certain parts of the door fall within the shadow of my desk. And yet, my experience of the area in the shadow is qualitatively different than that of the well-lit area. This is not a change in colour, but merely in lighting conditions. Kelly takes this to show that:

> **. . . the complete and accurate account of my perceptual experience of the colour of an object must contain some reference to the lighting context in which that colour is perceived.**[87]

Imagine, for instance, if I were presented with a patch of red for which I had no ready concept, so that I could only say "that shade of red" in pointing it out. Kelly is suggesting that this demonstrative concept could not specify the sort of lighting conditions under which the colour was seen; whether it was seen in the sun or in the shade, for instance. But since differences in lighting are nonetheless represented in the experience, it follows that the demonstrative concept is inadequate in accounting for that experience.

---

[87] Kelly (2001) p. 228

The case of a colour's dependency on the object in which it inheres makes a similar point. In a discussion of colour in his Phenomenology of Perception, Merleau-Ponty remarks on the object-dependency of colour, saying that "the blue of the carpet would not be the same blue were it not to be a wooly blue."[88] Kelly takes up Merleau-Ponty's point as an implicit argument against demonstrative concepts. A blue sphere made out of a shiny metal sits on a blue carpet. Objectively, the same shade of blue is shared by the surfaces of either object; we could measure the wavelengths of the light they reflect to satisfy ourselves of this. But despite that objectively they are the same colour, the two instance of "that shade of blue" appear phenomenally to be very different, so that the demonstrative I just used in the last clause of this sentence is not specific enough to capture either my experience of the ball or the carpet, though it might very well refer 'successfully' to the colour itself in abstraction from the experience.

The gripe Kelly seems to be making with both of these examples is that demonstrative concepts are, of necessity, general, but that every instance of perception will be a specific instantiation of the property to which that concept refers which is embedded in an environment and inheres in an object. But since the demonstrative concept is neutral between the character of each of these

[88] Merleau-Ponty (1945) p. 365

experiences – it merely points to that which they have in common – it fails to capture the total character of the experience:

> **A demonstrative concept like *that scarlet* can only pick out one scarlet among others. But the difference between the experience of the scarlet scarf and the experience of the scarlet steel ball is, ex hypothesi, not due to a difference of color (this shade of scarlet versus that shade of scarlet), but rather is due to a difference in the object that manifests that color.**[89]

Kelly believes that if these two observations about perception are accurate this would show, *contra* Peacocke, that "demonstrative concepts are too coarse-grained, not too fine-grained, to capture perceptual content."[90]

We have seen in this section that perceptual experience cannot be fully captured by demonstrative concepts since those concepts cut too coarsely to capture certain aspects of experience. I also gave some reasons for believing that an account of how we came to have conceptual content may rely on the prior existence of nonconceptual content. Contrary to McDowell's claim, it seems that sensation therefore *must* make a separable contribution to experience.

B. Nonconceptual content and "islands of rationality"

It seems to me that an appeal to nonconceptual content is most often seen by conceptualists to come dangerously close to taking up the 'sensation' view of

---

[89] Kelly (2001) p. 229

[90] Ibid. p. 229

perception once more. Conceptualists feel that appeals to sensations and nonconceptual content are inadequate for explaining the content of perception since perceiving, they contend, involves a level of interpretation of the perceptual field which necessarily employs the use of concepts. That perception is dependent on concepts is indicated, as Hamlyn claims, by the fact that "in order to perceive something we must perceive it *as such and such.*"[91] All perceiving is perceiving-as. Although he does not endorse this view, Peter Carruthers nicely encapsulates it when he says:

> **One might wonder whether [nonconceptualism] is even coherent. For there can surely be no *representation* with *discrimination* – you cannot *represent* red unless you can *discriminate* it from other colours – yet all discrimination might be thought to implicate concepts.**[92]

This squares with the conceptualist John McDowell's claim that there is no way to understand phenomenality – that is, how a perception can be "a seeming glimpse of the world" for the organism – apart from involving the deployment of conceptual capacities: so, for instance, McDowell remarks that:

> **. . . although Evans [McDowell's nonconceptualist interlocutor] does take care to credit experience with content, that does not save them from being intuitions in a sense that entitles us to apply the Kantian tag to them: since they are without concepts, they are blind.**[93]

---

[91] Hamlyn (1994) p. 91

[92] Carruthers (2000) p. 130

[93] McDowell (1996) p. 54

As I see it, there are two reasons why McDowell feels that only conceptual content can grant us sight. First, the organism must have some kind of self-conception – some *concept* of itself as the subject of perception – which allows it to grasp the self/world distinction which is a prerequisite for experience, and this requires an explicit concept of oneself as the subject of experiences, as an agent, etc. Secondly, the perceptual scene itself must contain real distinctions between things in the world and not simply a morass of undifferentiated informational states, and we can only distinguish one thing from another if we carve up the perceptual scene with concepts. Normally, or *naturally*, when I look at the world what I see are various *objects*, and not merely an array of undifferentiated sensations. But it seems obvious that identifying something as an object, and thereby delimiting it from other objects in the sensory array is a paradigmatically *conceptual* activity. It is McDowell's belief that nonconceptual content cannot fulfill either of these role.

I disagree. Although it is obvious that perception implies *some* manner of discrimination – whether of self or world – we do not have to go so far as to posit that these discriminations must always be conceptual in nature. If we were to insist on this – as conceptualists do – we would be glossing over a whole range of intermediate states of discrimination which present the world in some way or other but which *do not* thereby generalize so that what is discriminated can be

reified and pulled out of the context in which it is encountered. To do so implies the possession of further abilities which, while they may enhance the amount of information a creature can cull from its phenomenal experience, are not *necessary* for experience as such. This idea is summed up by Bermudez when he writes: "What leaves space open for the possibility of nonconceptual content, however, is that compositional structure can exist in the absence of global recombinability."[94]

I will now relate three examples which I believe clearly meet this criterion for nonconceptual content. In each case, there is a sort of discrimination – either of itself *qua* agent or the environment – which an intentional agent utilizes in deliberating on a course of action – a deliberation which has the mark of content in that it can go awry or misrepresent the situation – but that discrimination itself is not immediately available to the agent in such a way as to be globally recombinable with other content. In other words, such content fails to meet Evans' Generality Constraint, which states that in order for a content to be conceptual it must not be bound to certain contexts of deployment. True concepts are discrete components; even if I don't fully understand the meaning of a concept I am using it is still possible for me to recombine it with any other concepts I wish, even if the resulting thought is nonsensical. Similarly, if I am having a conceptual

[94] Bermudez (1998) p. 93

experience with the content that could be specified in thought, then it is within my ability to freely imagine the content of that experience being abstracted out from the current scene and placed in novel situations. *Nonconceptual* contents will therefore be discriminations which are made that are limited to certain "context-bound islands of practical rationality," to use Susan Hurley's phrase.[95]

So, for instance, assuming that content is that which is specified in the "that" clause of a statement, the nonconceptual experience enjoyed by a frog as a fly passes through its field of vision would be rendered: "The frog sees that there-is-a-fly-over-there." The hyphenation suggests that this content cannot be decomposed. To put it another way: the content of the frog's experience enters into rational relations with its immediate goals (of catching and eating the fly) while failing to generalize and thereby become part of a conceptual repertoire. The analogous experience in a human would be rendered in the normal way: "Jim sees that there is a fly over there." Each element in Jim's conceptualized experience of the fly is available to him in an explicit way and can therefore be abstracted out of this experience and considered on its own or in connection with other elements remote from the immediate situation: flies cannot just be seen "over there," but they can be seen elsewhere (or in paintings), considered to be

---

[95] Hurley (1998) p. 138

bothersome, believed to carry disease, etc. Jim, unlike the frog, can think about flies in any context he wishes.

The upshot of all of this is twofold: first, that there can be nonhuman experience without conceptualization; second, that one can have personal-level reasons without those reasons being part of an epistemological project of justifying beliefs – some reasons are reasons for *acting*, and not reasons for *believing*.

First, drawing on the work of Jose Luis Bermudez, I will argue that we can make sense of a first-person point of view without resorting to concepts. Second, I will look at an example from Andy Clark , who suggests that the "knowledge" of a non-linguistic biological creature is best described as deriving from mental states bearing nonconceptual content. My third example is from Adrian Cussins, who illustrates the relation between situated, real-time action and nonconceptual content through an analysis of the sort of knowledge involved in the skilled operation of a motorcycle.

i. Nonconceptual points of view – McDowell notes that Gareth Evans, a nonconceptualist, admits awareness only takes place:

> **. . . against the background of an understanding of how perception and reality are related, something sufficient to sustain the idea that the world reveals itself to a perceiving subject in different regions and aspects, in a**

way that depends on the subject's movement through the world.[96]

In other words, awareness is dependent upon the 'background' possession of a point of view, and of an understanding of how that point of view is related to the world. But, McDowell asks, what in turn makes this "background" possible? The short answer is "a *concept* of self." Awareness of the world presupposes that some manner of distinction has already been made between self and world, and this separation is made by having some (possibly rudimentary) *concept* of one's self: "Having things appear to one a certain way is already itself a mode of actual operations of conceptual capacities."[97] Minimally, having an experience presupposes that the organism comprehends (whether implicitly or explicitly) the relation between their perceptions and the objective reality which their perceptions are about. Alva Noë, for instance, puts the matter this way: "to be a perceiver, one must be capable of keeping track of the ways in which one's perceptual experience depends on what one does, and also more generally on one's relation to the world around one."[98] But this is as much as to say that an un-conceptualized experience is impossible.

This answer, as McDowell himself points out, has much in common with

---

[96] Ibid. p. 54

[97] McDowell (1996) p. 62

[98] Noë (2002) p. 185

Kant's claim in the Critique:

> **The *I think* must *be able* to accompany all my representations; for otherwise something would be represented in me that could not be thought at all, which is as much as to say that the representation would either be impossible or else at least would be nothing for me.**[99]

Both Kant and McDowell agree that the possession of this 'I think' is a conceptual accomplishment; note that McDowell refers to the subjective pole of experience as a necessarily containing "self-conscious *conception* of how [the subject] relates to the world," and that we cannot make sense of the possession of a self in the absence of concepts:

> **It is spontaneity of the understanding, the power of conceptual thinking, that brings both the world and the self into view. Creatures without conceptual capacities lack self consciousness and – this is part of the same package – experience of an objective reality.**[100]

But I think we can make sense of what it is to possess a point of view on the world and an understanding of how one's experiences relate to reality without invoking a *concept* of self. Jose Luis Bermudez has developed a theory in which nonconceptual forms of self-awareness are "logically and ontogenetically more primitive than the higher forms of self-consciousness that are usually the focus of philosophical debate."[101] The higher forms of self-awareness which employ

---

[99] [B 131-132]

[100] McDowell (1996) p. 114

[101] Bermudez (1998) p. 274

concepts, such as the explicit awareness of oneself as the subject of experience, are built upon a strata of simpler nonconceptual forms of self-awareness. To secure my point I will briefly explain Bermudez' account of one form of nonconceptual awareness.

Bermudez' analysis draws on the notion of an *affordance* which is found in the work of J.J. Gibson. An affordance is a higher-order invariant in perception through which a creature detects properties of its environment that are relevant to its abilities (such as walking, climbing, eating, etc.) The perception of an affordance is relative to the abilities of the animal perceiving it; so, for instance, an ape will see a branch as *affording* the possibility of brachiation, while a bird might perceive it as *affording* the possibility of perching. But this also means that the affordance cannot be reduced to the physical properties of the object considered in abstraction; they are necessarily defined in relation to the acting creature. The brittleness of a branch might make it such that it doesn't afford the possibility of brachiation, but does afford the possibility of perching, for instance. We might be tempted to think that an affordance might be detected though some conceptual scheme possessed by the perceiving creature, but Gibson denied this. Rather than being specified through any sort of inference, an affordance is a property of the environment itself and, as such, is relative to the creature's set of possible actions rather than its set of mastered concepts.

What Bermudez means by saying that the perceptual environment contains "self-specifying" information is that it contains such affordances for the creature; in "the perception of affordances is a form of self-perception, or at least a way in which self-specifying information is perceived," since the "whole notion of an affordance is that of environmental information about one's own possibilities for action and reaction."[102] What is important for our concerns is Bermudez' claim that informational-perceptual states can contain "self-specifying information" that falls short of conceptual self-awareness; put differently, creatures with no conceptual self-awareness can nonetheless perceive the environment as containing affordance for *their* action.

To give one example, we can see this from an analysis of reaching behaviour in infants. Reaching behaviour, as Bermudez points out, "is driven by the perception that an object is within reach (by the perception that the object affords reaching), and this is, of course, a form of self-specifying information."[103] In the case of even very young infants there is experimental evidence of the perception of affordances, since the child will adjust its reaching behaviour in relation to the distance of the object, such that a behavioural difference is

[102] Bermudez (1998) p. 113

[103] Bermudez (1998) p. 127

observed between an infant's attitude toward an object which is within their reach and an object which falls outside their reach. Further, in children as young as two weeks old "when an object is placed out of all possible reach, there is a marked reduction in the frequency with which infants extend their arms toward it.[104] It is hard to explain the results of these experiments without accepting that the infant responds in these various ways because the infant is aware of the position of the object in relation to *itself*. What we should take from this is that the ability to accurately assess from one's perception whether or not an object is in reach develops in early infancy, that is, in humans who clearly have not mastered the use of conceptually based "I" statements; the clearest way to account for this ability is to accept that the child possesses some nonconceptual sense of self which informs its attitudes toward affordances.

ii. Skilled performers and creative cognizers – In "Connectionism and Cognitive Flexibility," Andy Clark makes a distinction between "skilled performers" and "creative cognizers." A creature which is merely a skilled performer is able to act in complex and flexible ways in a restricted domain. Clark's example is the beaver, which is able to build dams of great complexity in ways subtly fitted to their local environment. However, a beaver fails to be a

---

[104] Ibid. p. 127

"creative cognizer" since this dam-building procedure itself cannot be taken as an object for the beaver's own cognitive processes: "It cannot operate on and amend the procedure at will, despite its plasticity relative to the local context."[105] Humans, on the other hand, are creative cognizers: "We are able to bring arbitrary elements of our knowledge to bear on many of the tasks we perform. We are not bound in respect of a given task, to any single conceptual space."[106]

Clark goes on to argue that what determines whether a creature (or AI, for that matter) is merely a skilled performer or a creative cognizer is the sort of content which is made available to the creature via its mental representations: a mere skilled performer will only have access to nonconceptual content, whereas a creative cognizer will have access to conceptual content. The beaver navigates skillfully through a domain which we would describe using concepts like "branch" and "water." But the beaver need not possess contentful mental states containing these concepts in order to perform tasks in this domain. He only need possess those contents which aid him in finding his way around the domain space. Moreover, the beaver's inability to have thoughts *about* the dam-building domain suggests that his experiences of things in that domain are unstructured relative to

[105] Clark (1994), p. 165

[106] Ibid. p. 165

other of the beaver's task domains, and that they therefore do not satisfy Evans' Generality Constrain for concept possession. Conceptual content, by contrast, is "structured content in which each element implicated in the specification of the thought has a separate significance for the creature, and can enter freely into combinations with elements of the creature's other thoughts."[107]

Some reason for believing that the content of the beaver's mental states fail to meet the Generality Constraint is that he will build a dam *wherever* he hears rushing water. Therefore, this behaviour can be induced by the sounds of rushing water emanating from a loudspeaker; though there is no water to stop up in the area, the beaver is unable to use this information to modify his behaviour. He cannot halt his futile dam-building behaviour by reasoning that branches should be used to build a dam only in a location where there is, in fact, water. This sort of "intelligent self de-bugging," according to Clark, would require *explicit* representations of the things in the task domain, that is, a higher level analysis which could only be made available to the beaver by way of concepts:

In short, the possession of abstract categories and appropriate control structures would enable the system [in this case, the beaver] to *isolate* the cause of the trouble and take *focused* action to resolve it.[108]

---

[107] Ibid. p. 173

[108] Ibid. p. 169

The important point for our consideration is that action which is guided by nonconceptual content can still be considered as actions undertaken for "reasons." Of course, the reasons for which an animal acts play no role in an epistemological project; animals which possess "islands of instrumental rationality" do not use these reasons to justify beliefs. But the reasons for which they act are nonetheless their *own* reasons, and not simply the results of sub-personal information processing.

To specify the content of phenomenal perception in conceptual terms is thus necessarily to fail to capture it: "ties between conceptual contents and actions are contingent whereas what it *is* to be in the nonconceptually contentful state just *is* to be disposed to move and act in certain ways."[109] Clark urges us instead to think of the structures which are made available to agents through nonconceptual content as "the epistemological bedrock which puts the system in *contact* with the world its thoughts are meant to concern."[110]

iii. Skilled, situated action and context-bound reasoning – I have been suggesting that animal behaviour and the behaviour of pre-linguistic infants present problems which are best accounted for by appealing to states with

[109] Ibid. p. 172

[110] Ibid. p. 166

nonconceptual content. Following Adrian Cussins and Susan Hurley, I would like to suggest further that nonconceptual content can also be used to explain the behaviour and knowledge of language-using adults. While it remains an open question – one which I will not attempt to answer here – as to whether conscious perceptual states in language use are *entirely* nonconceptual, I think these examples establish that there is a useful way in which these states can be described as at least *partially* nonconceptual.

Cussins considers two way in which he might know about the speed he is travelling on his motorcycle to illustrate this. He relates a story about being stopped by a police officer for speeding while riding his motorcycle through the streets of London, and being asked whether or not he knew how fast he was going. Though he had not looked at the speedometer for information on his speed, Cussins suggests that there is still a sense in which he *knows* how fast he is travelling. This is because he had been utilizing some knowledge of his speed in the course of "knowingly making micro-adjustments of [his] speed all the time in response to changing road conditions."[111] Further, he claims that he is as epistemically responsible for those micro-adjustments as he is for his normal judgements, since the knowledge which he possesses of his speed provides him

[111] Cussins (1990) p. 150

with *reasons* for undertaking certain actions such as slowing down, speeding up, leaning just so, wiggling through dense traffic, etc. However, this knowledge of speed is bound in important respects to the situated action of controlling the motorcycle: "This kind of knowledge of speed does not entail that I be able to recognize it as the same speed again as I rode down an uncluttered motorway outside the city."[112]

This is because Cussins did not know the speed he is travelling *as a speed.* If asked, he could not say whether he was travelling at 50 mph, for instance. He simply did not possess the right sort of information – a *conceptual* knowledge of the speed at which he was travelling – to make that statement in an epistemically responsible way: "the speed of my motorcycle was not made available to me as that which would render true certain propositions, and false certain others."[113] Had he known his speed *as a speed* – as an "object," in Cussins' manner of speaking – he would have been able to use this knowledge to make inferences outside of this situated, action-oriented domain. For instance, he could present it to himself, to the police officer, to the traffic court, etc. as the same speed in all of these different contexts. Because of this, Cussins suggests that the experience

[112] Ibid. p. 150

[113] Ibid. p. 150

which he had of his speed while he drove through the streets was nonconceptual:

> **. . . this kind of content cannot by itself provide what we have come to regard as the constitutive requirements on *thought* content: generality, objectivity, standardization, transportability of knowledge from one embodied and environmentally specific situation to another.**[114]

Susan Hurley makes a similar point regarding the Wason effect, but in this case inferential skills are bound to *social* contexts instead of contexts involving situated action. However, the same results seem to follow: an agent can act for her own reasons even though her ability to act upon those reasons does not generalize to other contexts. Hurley relates that Wason asked people to test a simple instance of "p implies q":

> **... if a card has "D" on one side, it has "3" on the other side. Subjects observed 4 cards, showing on their upturned sides: D, F, 3, 7. They were asked which cards they should turn over to determine whether the rule was correct. The right answer is: the D card and the 7 card. Most people (90-95%, including those trained in logic) choose either just the D card or the D card and the 3 card.**[115]

But despite their poor performance in performing this inference with these symbols, people fared much better when tested on a instance of "p implies q" that described an exchange of the form: if you take a benefit, you must meet a requirement. Hurley suggests that performance improved because people are very

---

[114] Ibid. p. 151

[115] Hurley (2001) p. 425

good at detecting cheaters in social situations. Imagine for instance that the exchange was rendered concretely as: "If a person is drinking alcohol then that person must be over 19." The error that was made in flipping over the "3" card in the abstract experiment seems glaringly obvious in this context. It would be the equivalent of trying to discover whether someone who was over 19 was drinking alcohol; this would be unhelpful in determining whether or not the rule was being broken. This result suggests that "when an agent acts on her perceptions so as to flush out a cheater, she can be acting on her own reasons, available from her point of view, even though they are not inferentially promiscuous."[116]

2.3 Summation

I have provided several examples of contentful states which fail Evans' Generality Constraint. The existence of contentful states which are not conceptual suggests that conscious creatures, qua conscious creature, enjoy nonconceptual yet structured representations of their immediate environment which help them in navigating through the world. Further, this representation manifests at the personal level; these actions we take through navigating the world are intentional and norm-governed – for instance, actions can be performed skillfully, or un-

[116] Ibid. p. 425

skillfully – but they are not undertaken within the epistemological project of determining the truth or falsity of our thoughts, of having consistent beliefs, and so on.

But I do not believe that conceptualism is simply a failed theory. Instead, conceptualism is perhaps best understood as a theory of how concept possession (and, more generally, language use) affords us the ability to "pay attention" and to parse the existing phenomenal content in ways unavailable to non-linguistic creatures, *rather than* (as conceptualists would contend) a theory about phenomenal content *simpliciter*. Conceptual abilities might afford the creature some higher-order thoughts about their perceptual states ("That's a cat!"), whereby they could reason *about* their activities in an explicit way, but they do not grant the creature its first glimpse of the world. Conceptualists simply overestimate the explanatory power of their theory.

I think this is enough to show that the perception of objects is not a conceptual affair in every respect. But we must not be led from the conclusion that certain aspects of experience are best explained by invoking a level of nonconceptual content to the further conclusion that the look of experience is fixed *entirely* by nonconceptual content. Nonconceptualism *simpliciter* would require additional arguments, which will I will consider in the next chapter.

3. Nonconceptualism

"I thought without words, *on* things, *with* things . . ."

- Jean Paul Sartre, *Nausea*[117]

3.1 Elucidation

In the last chapter conceptualism was introduced as the modern progeny of a longstanding philosophical tradition which sought to legitimate a distinction between perception and sensation. An organism's dealings with the world, goes the theory, are conceptually structured perceptions, with sensations playing only an ancillary and purely causal role which is of no consequence to the creature itself. Put another way, conceptualism is the view that the ways in which creatures can represent the world are constrained by their conceptual capacities. Sensations for a conceptualist are non-intentional mental states, while perception, which certainly *is* intentional, is always a conceptual affair.

But this doctrine didn't seem to square with several observations about the character of perception. The richness and fineness of grain that perceptual experience exhibits, and the unmediated, "non-doxastic" ties that it has with action, suggest that perceptual experience is not exhausted by the concepts that the perceiver possesses. For these reasons, in the last chapter I endorsed *partial*

[117] Sartre (1938) p. 129

nonconceptualism, which I understand to be the view that sensations *do* have intentional content, though that content is *often*, though perhaps not always, different from the content of thoughts. Note that this is not a simple reversion to the causal-sensational view of Aristotle et al. Whereas 'sensation' theories of perception understood that process as a mere causal impingement on the sense organs from the external world, nonconceptualism is an attempt to understand sensations as *full-blooded representations*. Sensations therefore will bear content or be 'about' parts of the world, have correctness conditions, and so forth.

I concluded that our phenomenology isn't made up *exclusively* of mental states with conceptual content. But nonconceptualists go further than this. Their claim is that phenomenal states are *always* nonconceptual, and that conceptualization (in the form of thinking) is always exterior to phenomenology. In this spirit Fred Dretske, an eminent defender of nonconceptualism, criticizes theories of perception such as conceptualism, which he says:

> **. . . conflate perceptual and sensory phenomena on the one hand with cognitive and conceptual phenomena on the other" and are therefore based on an erroneous assumption that seeing and hearing are "low-grade forms of knowing.**[118]

Such a theory, in Dretske's words, could only serve to "obscure the distinctive role of *sensory experience* in the entire cognitive process."[119]

---

[118] Dretske (1981) p. 135

[119] Ibid. p. 135

My discussion of nonconceptualism will mostly centre on the theory of phenomenal perception developed by Dretske in his book Naturalizing the Mind. Certain particularities of Dretske's theory are not shared by most other nonconceptualists; where necessary, I will supplement his theory with the work of other theorists to render a more balanced view of nonconceptualism As I did in the last chapter in dealing with conceptualism, I will summarize the arguments and motivations for adopting nonconceptualism before moving onto some critiques of the theory.

A. Indicator functions

Though I have already given a rough picture of representationalism in the first chapter it will be useful here to briefly relate Dretske's overview of the theory, since he carries the vocabulary established in this overview into the rest of his argument. "The fundamental idea [behind "representation"]," says Dretske, "is that a system, S represents a property, F, if and only if S has the function of indicating (providing information about) the F of a certain domain of objects."[120] In order to do this, a representational system must have states (s_1, s_2, . . . s_n) that correspond (or, perhaps causally co-vary) with different values (f_1, f_2, . . . f_n) of a property. The "ready" light on a stove is an example of a very simple representational system with only two states: the light is on until the oven reaches

[120] Dretske (1995) p. 2

the desired temperature, at which time it switches off. A more complex "ready" light might represent gradients of heat by slowly dimming as the oven heated. Or we could get rid of the light and instead have a digital display which numerically represented the oven's temperature. Each of these "indicators" is part of a representational system which has the function of providing me with information about the temperature inside the oven (though some representational systems are better at discriminating states than others). And that they have this function implies that it is possible they might fail to fulfill that function, causing the systems to *misrepresent* the state of the property it is tracking; for instance, if the light was improperly wired it might get brighter instead of dimmer as the temperature in the oven rose. Merely causally co-varying with a property is therefore not enough to make a system representational; as Dretske points out, the path that drifting smoke takes in rising through the air may be correlated with wind speed, but since it does not have the function of representing this property it could not possibly *mis*represent its value in the way an anemometer can. Likewise, a black-and-white TV does not misrepresent the colour of the sky since it does not have the function of supplying information about colour. But a colour TV displaying a black and white image of the sky does misrepresent the sky's colour.[121]

We have so far been talking about representational systems which acquire

---

[121] Ibid. p. 4-5

their functions from the intentions of their human creators. Man-made artifacts acquire their function by *convention*, but this is not so in the case of mental states. All mental states, conceptual and nonconceptual alike, have the naturally (that is, evolutionarily) acquired functions of indicating the states of properties in certain domains of our environment (both distal and internal). It is not simply a bit of folk psychology that we describe sensory organs in terms of what they are "for": "The senses yield representations of the world, not just because they (when working right) deliver information about the world, but because that is their job."[122] Mental states are therefore *natural* – not conventional or man-made – representations.

B. Two senses of "looks"

Dretske makes a further distinction between the way mental states represent. The visual experience of Paul playing the piano is a different kind of representation than the *belief* that he is playing the piano. Visual and auditory experiences of piano playing do not require concepts, or any sort of understanding of what a piano is or sounds like, so that even those creatures which entirely lack concepts can hear (that is, experience) pianos being played. Having the belief that a piano is being played, on the other hand, requires the concept 'piano'.

Some evidence that beliefs – and therefore concepts – are not necessary for having experiences is that these two sorts of representations can come apart. One

[122] Ibid. p. 5

can hear a piano without believing that one is being played. A dog, for example, can certainly hear a piano though he lacks the concept 'piano.' Conversely, one can believe a piano is being played without seeing or hearing a piano.

Optical illusions are another example of this difference. Knowing that one is subject to an optical illusion does not eliminate that illusion, as we might expect if experiences were belief-dependent: "In cases like these it makes sense to distinguish between perception and (perceptual) belief, and to identify a non-doxastic (non-belief) component in perceptual experience."[123]

According to Dretske, we must take care when describing experiences so that we do not jumble them up with a belief-experience hybrid such as an awareness or consciousness; these words often imply a conceptual representation – a *thought* – that does not necessarily occur alongside the experience. For instance, we could imagine that:

> **[a] child or an animal might be visually aware of the shirt's color (their visual experience of the shirt's being, as they say, suffused with blueness) without their knowing or thinking that the shirt is blue – without sorting (or having any disposition to sort) the shirt with other blue objects.**[124]

The experience of the shirt's colour, Dretske suggests, does not require an attendant conceptual awareness; I can see a blue shirt without seeing it *as* a blue shirt.

---

[123] Bermudez (1995) p. 185

[124] Dretske (1995) p. 11

Noting that many "appear" and "look" words apply ambiguously to both experiences and belief-experience hybrids, Dretske coins a bit of terminology to keep matters straight. He asks us to imagine two scenarios: Susan – "a child of normal eyesight and intelligence" – has never seen nor heard of dogs. The first dog she sees is a French poodle. Granted that Susan will not *say* or *think* that what she is seeing looks like a poodle, will it nonetheless *look* like a poodle to her?; Arthur – "a toad of normal toad-eyesight and toad-intelligence" – sees the same poodle. Like Susan, he has no beliefs about dogs, and does not conceptualize his experience as that of a poodle. But is there a sense in which it looks like a poodle to him?"[125]

i. Looks$_p$ – In one sense of "look," we will want to say that it does indeed look to Susan like there is poodle in front of her. As a human with normal eyesight, she is able to visually discriminate poodles from other breeds of dogs (and other medium sized objects) in the same way as any other human with normal eyesight, even if she cannot describe, or conceptualize, the way the dog looks to her as "like a poodle."[126] But that she has this cognitive deficit, Dretske suggests, does not make a difference in how the dog *looks* to her. Dretske calls this the *phenomenal* sense of look: "To say that a dog looks – phenomenally – like a poodle to S (=looks$_p$) is to say two things: (1) that the dog looks to S the

[125] Ibid. p. 66

[126] Ibid. p. 66

way poodles normally looks to S; and (2) the dog looks different to S from other dogs (bulldogs, terriers, etc.)"[127]

Arthur (the aforementioned toad) does not possess the visual acuity of humans with normal eyesight. Citing behavioural data, Dretske claims that Arthur's *normal* vision would give him the visual acuity of a half-blind human. Notice the difference: though the visual phenomenology of a half-blind human and a toad while looking at a poodle may be indistinguishable, the former is a misrepresentation while the latter is not (much like the black-and-white TV example above). This is because Arthur's visual system, unlike the half-blind human, doesn't have the natural *function* of discriminating poodle-sized objects. Arthur will not see the differences of form and detail that make poodles look$_p$ different than bulldogs to humans with normal eyesight, but this is as it should be. For Arthur, poodles will *naturally* look like bulldogs, which will look very much like fire hydrants, etc., that is, they will all look something like blurry spots.

ii. Looks$_d$ – In another sense of "looks," we will want to say that it does not look to Susan (nor to Arthur, for that matter) *as if* there were a poodle in front of her. Since she lacks a command of the concept POODLE, her perception of the dog does not cause her to *believe* that there is a poodle in front of her: she would not classify, identify, or report her experience as that of a poodle. Susan therefore lacks what Dretske calls the *doxastic* sense of "look." This sense of "look" words

---

[127] Ibid. p. 68

has close ties to beliefs: "To say that a dog $looks_d$ like a poodle is to say that, in the absence of countervailing considerations, this is what S would take the dog to be, what S's perception of the dog would (normally) prompt her to believe."[128] Again, it is Dretske's claim that an object $looking_d$ some way or other will not alter how it $looks_p$.

C. Two senses of "represents"

It may already be obvious that $looks_d$ and $looks_p$ map onto thoughts and experiences, respectively. But how can we draw a principled distinction between thoughts and experiences, between $looks_d$ and $looks_p$, in representational terms? Dretske suggests that the difference in these representations – on the one hand, the experience of poodles, etc. and, on the other the belief about poodles, etc. – stems from differing sources of their indicator functions.

i. Represents$_s$ – Perceptual systems, as has already been mentioned, are described in terms of what they are "for." Dretske points out that this is a representational way of thinking about perceptual systems. So for instance:

> **. . . semicircular canals of the middle ear are said to *be for* the detection of indication of linear acceleration . . . and the retina for encoding information about light for transmission to the brain.**[129]

The indicator function of experiences are therefore produced by the natural function of the perceptual systems in which they are contained. The natural

---

[128] Ibid. p. 68

[129] Dretske (1995) p. 5

function of perceptual systems, in turn, is to provide some information about either the distal or internal environment. Certain visual systems have the systemic function, which Dretske shorthands function$_s$, of delivering information (via representations$_s$) about colour. It is unimportant for the point Dretske is making that there is some debate over which properties of the distal environment colours actually *are*. He is simply insisting that an experiential state of, say, red, has the function$_s$ of indicating the presence of red; that is, the state represents$_s$ some feature of the world as being red.

Representations$_s$ – the quality of sensory states at their "most basic (phenomenal) level" – are products of phylogenetic development: "experiences have their representational content fixed by the biological functions of the sensory systems of which they are states."[130] In other words, the qualities of experience *just are* the properties which objects are represented$_s$ as having. To return to the above example, our phenomenal experience of poodles differs from that of toads because "[a] toad's visual system does not represent$_s$ poodles the way our visual system represents$_s$ them."[131]

ii. Represents$_a$ – On the other hand, a *belief state* does not acquire its indicator function from the system of which it is a state, but "from the *type* of state

---

[130] Ibid. p. 15

[131] Ibid. p. 70

of which it is a token."[132] In other words, a representational$_s$ state of our sensory system may "acquire" an additional indicator function$_a$. These representations$_a$ are not "built-in" to the system An experiential state of red *"k"*, for instance, might also come to *mean* "red" or, in certain contexts, it may *mean* "stop." But these additional indicator functions won't do much to change the nature of the initial sensory representation: "Through learning, I can change what I believe when I see *k*, but I can't much change the way *k* looks (phenomenally) to me, the kind of visual experience *k* produces in me."[133]

Dretske imagines two of Pavlov's dogs as an example of the distinction he is making between mental representations. One dog has been conditioned to salivate whenever it hears a middle C being played, while the other has been conditioned to salivate when it hears a clarinet being played. Both dogs hear the same sound when a middle C is played on a clarinet; they presumably have identical experiences., and they respond the same to the stimuli by salivating. But, Dretske says, their responses are mediated by different representations.

> **As a result of different learning, the dogs hear it differently – one ... hears it as middle C, the other as the sound of a clarinet. The way their experience represents$_s$ the sound may well be the same, but the way their experience represents$_a$ it is different.**[134]

---

[132] Ibid. p. 13

[133] Ibid. p. 15

[134] Ibid. p. 15

We could say that the dogs *label* their shared experience using two different beliefs, neither of which can capture the richness or significantly alter the content of the experience upon which they are based.

There is a related way of drawing the distinction between experiences and beliefs that is favoured by Michael Tye. For Tye, perceptual sensations:

> **. . . form the outputs of specialized sensory modules and stand ready to produce conceptual responses via the action of higher-level cognitive processing of one sort or another.**[135]

When sensory information is in this position Tye says that it is "poised" to make an impact on cognition. Using Dretske's terminology, we could say that the representation which stands as input to the cognitive processes is a representation$_s$. The same state, once it enters into the cognitive system, can also become a representation$_a$: experiences "are those natural representations$_s$ that service the construction of representations$_a$, representations$_s$ that can be calibrated (by learning) to more effectively service an organism's needs and desires."[136]

D. Fineness of grain

If we accept the above picture, it follows that phenomenal consciousness (something looking$_p$) is nonconceptual. Thoughts, on the other hand, seem to be limited by the range of the concepts which the creature possesses. Forming the perceptual belief that something looks like a cat, for example, necessarily involves

---

[135] Tye (1995) p. 103-104

[136] Dretske (1995) p. 19

remembering what cats look like. On the standard model of perceptual memory, perceptual information about what cats look like is stored in a schema (which, for our purposes, we can understand to be identical to the concept "cat").[137] Once the sensory input is brought under this schema, the subject can form the belief "That is a cat." Perceptual beliefs, then, as Tye points out, are constrained by limitations on memory: "If one lacks appropriate schema, the corresponding belief cannot occur."[138] But the features which enter into the content of phenomenal content – those representations$_s$ produced by our perceptual systems – are not limited in this way to those features for which the subject possesses concepts. Something cannot look$_d$ red if I do not possess the concept "red," but surely it can still look$_p$ red to me.

Another way to say this is that perceptual states (something looking$_p$) are more *fine-grained* than any concepts which we might have for them. By fine-grained, I mean, for instance: the complex shapes drawn in the air by cigarette smoke, or the many unnamed shades of red we are able to discriminate in colour vision; the complex phonological signatures of speech; and the constantly varying sensations we have while running a hand over wool cloth. We could even make a case for the perceptions of spatial distances being fine-grained. I might not *explicitly* know that I am standing exactly 200 metre away from University Hall

[137] Tye (1995) p. 66

[138] Tye (1995) p. 66

and 80 metres from Mills Library, but I nonetheless easily understand that my distance from University Hall is of a greater (though, for me, indeterminate) magnitude than my distance from Mills Library. In each of these cases, we are presented with aspects of experience for which we seem to have no readily available concepts. The concepts "smoke" and "red" and "wool," and the linguistic meaning of the utterances in speech – none of these concepts seem to do full justice to these experiences. They are simply too crude. Even more specific concepts in these cases – crimson, superfine merino – are again too coarse to fully capture the fine content of these representations. The case of spatial perception is a slightly different case: although we possess the concept metre it isn't necessary that we employ this concept in specifying the content of my perception. I can know something (admittedly more vague) about my position relative to University Hall and Mills Library without employing any concepts that would specify this position.

Let's take a look at a paradigmatic case of colour perception as a concrete example of how our perceptual experiences might be more fine-grained than our conceptual repertoire. It seems obvious that we are able to discriminate far more colours than we have names for. If I had time to think about it, I might come up with a dozen words I know for different shades of red, though I would definitely be shaky with matching these terms to actual shades. But I can nonetheless discriminate countless of different colours despite my relative paucity of colour

terms, which suggests that my experience of colour is not mediated by concepts. For example, I will *experience* red_{32} and red_{27} as two distinct shades, even though I might have no fixed concepts for either. If I were to look at a chart of different shades of red, I would not be able to *indicate* which among them was red_{27} – in other words, I could not apply the concept red_{27} – even though I nonetheless *experience* it as being distinct from other reds. Or consider if I was shown a sample of red_{27} and then, a short time later shown it again next to a sample of red_{28}. It is likely that I would be unable to say which was the red_{27}. This seems to be compelling evidence for at least partial nonconceptualism in the content of experience.

No one involved in the debate wants to deny that experience contains fine-grained content; it is plainly obvious that it does. But granted that the richness of experience is part of the content of representational states, what sort of content are they? The so-called "ineffability" of fine-grained experience certainly seems to square with the intuition that fineness of grain is best accounted for by nonconceptual content. After all, nonconceptual content cannot be adequately described conceptually, so we will not do justice to phenomenology:

> **. . . if we restrict ourselves to those contents which can be built up by referring to the properties and relations which the perceived objects are represented by the experience as possessing.**[139]

There are some responses open to conceptualists, though I do not think

[139] Peacocke (2001), p. 3

that they work. For instance, McDowell responds by suggesting that we should not limit conceptual thinking about colour to concepts expressible by definite colour names such as "red," "ochre," etc. We can also *acquire* as many colour concepts as there are colours, so that one's concepts can capture colours no less finely than experience presents them. The gaps between our specific colour concepts can be filled, says McDowell, with demonstrative concepts:

> **In the throes of an experience of the kind that putatively transcends one's conceptual powers – an experience that *ex hypothesi* affords a suitable sample – one can give linguistic expression to a concept that is as exactly as fine-grained as the experience, by uttering a phrase like "that shade," in which the demonstrative exploits the presence of the sample.**[140]

It is important that this demonstrative concept, although formed from an instance in which the colour sample is actually perceived, can be used beyond the duration of that colour experience. Otherwise, its status as a concept would be suspect since it would lack the Principle of Compositionality, which requires that in order for something to count as a concept we must be able to utilize it in novel situations. In the case of demonstrative colour concepts, McDowell concedes that our opportunity for using these concepts might only be a small window of time directly after the act of perception in which it was formed, after which it would presumably fade out of our conceptual repertoire. But part of what makes it a concept is that we *could*, for instance, remark after the colour sample was taken away: "That shade looked a lot like the colour of the drapes in my old apartment!"

---

[140] McDowell (1996) p. 56-57

This would be an instance of a demonstrative colour concept entering into rational relations with other concepts. And this is enough to ensure that the initial perception it was culled from was conceptualized, though not *explicitly*, since although the experience doesn't give rise to a definite colour concept in us we are still able to mobilize "recognitional capacities" in identifying the shade.[141] If McDowell is right in characterizing this as an instance of possessing a concept of the unnamed shade then the nonconceptualist's claim that fineness of perceptual grain necessitates a nonconceptual level of content will have failed.

In an attempt to head off criticisms, McDowell also considers the possibility that these demonstrative concepts might be "hybrids" which depend on both nonconceptual "intuition" and conceptual "understanding" for their existence. But this, he says, would presuppose the dualism between concepts and intuition without actually arguing for it. I think that he is wrong about this; far from simply waving our hands, we can give good reasons for believing this. Christopher Peacocke develops an account of how demonstratives are dependent upon the ways in which we experience things. He also argues, contra McDowell, that these demonstratives cut *too finely* in order to capture the experiences to which they refer.

Peacocke distinguishes three levels of description that are applicable when perceiving the shape of an object (though the choice of 'shape' as opposed to

---

[141] Ibid. p. 58

'colour' or 'smell' or 'taste' is of course arbitrary):

(i) The particular shape itself.
(ii) The way the shape is given in experience.
(iii) The demonstrative concepts which are made available by the way the shape is experienced.[142]

Peacocke notes that the nonconceptualist will say that items in (ii) and (iii) will stand in a many-one relation to those items in the immediately preceding levels. So, for instance, there will be more than one way of seeing a shape; a square with one of its sides resting on a plane will be seen *as a square*, but if it is balancing on one of its corners it will most naturally be seen *as a diamond.* And there will be more than one demonstrative concept for identifying the experience of the diamond: "that diamond," "that shape," or "that pointy thing," to name a few.

McDowell's point about demonstrative concepts is that the ways featuring in (ii) will be entirely specifiable via demonstrative concepts in level (iii); in other words, every experience will have a demonstrative concept with which it can be captured. But according to Peacocke demonstrative concepts actually "slice *too* finely to capture the ways of level (ii)"; they overshoot their mark.[143] For which of the above descriptions of the shape with its point on the plane will be the one which correctly captures the experience – "that diamond," or perhaps "that pointed thing" – and further, what reason could we have for claiming that one is

---

[142] Peacocke (1998) p. 381

[143] Ibid. p. 382

more appropriate than another?

One possible response might be that the *most specific* demonstrative will be the one which directly captures the experience. But this would imply that two observers who had different "most specific" demonstrative concepts would be having distinct experiences at the most fine-grained level of experience. Imagine two observers that lack the general concept of "shade," one of whom has the concept "crimson" and the other "scarlet." On McDowell's view, when presented with a red colour patch they would have two different experiences.

Another possible rebuttal that conceptualist might make is that the general concept employed is irrelevant so long as the same *referent* is being picked out by one individual's "that scarlet" and the other's "that crimson"; that is to say, each phrase *picks out* the exact same property inhering in the world and therefore exactly picks out experience of that property. However, since the referent is a property of the *world* and not the way in which one *experiences* the world an appeal to the referent wouldn't capture (ii), but the actual shape itself at (i). And clearly we cannot get from the fact that the same actually existing shape, color, tone, etc. is being experienced to the further claim that it is being experienced in the same *way* by both individuals. So for example, when we hear middle-C and F-sharp played together on a piano, we might hear this interval as an augmented fourth, or as a diminished fifth. To use the demonstrative "that tone" will

therefore not capture the character of either experience.[144] Contrary to McDowell's claim, it seems that sensation therefore *must* make a separable contribution to experience, since not even general demonstrative concepts can fully specify the content of an experience.[145]

In summary, nonconceptualism consists of several core ideas: **(1)** that perceptual experience has a non-doxastic and concept-free component that is nonetheless representational; **(2)** that this concept-free representation is richer or more fine-grained than can be fully specified through concepts; **(3)** that phenomenal mental states *just are* these fine-grained perceptual experiences produced by the sensory systems and poised to make an impact on the cognitive system. **(3)** in turn entails two commitments of nonconceptualism that I would like to call into question: **(a)** that what we might call "cognitive perception" (something looking$_d$, or seeing something *as an "x"*) involving concepts does not alter the phenomenal content of experiences; **(b)** that only nonconceptual

---

[144] Peacocke (2001) p. 241

[145] It is not incoherent to suggest that demonstrative concepts are too coarse in dealing with certain aspects of experience, and too finely individuated to deal with others. Peacocke is saying that there are too many demonstrative concepts we might use to capture an experience, and that we cannot give any criteria for believing one of the demonstratives to be better than the other; in that sense, they are too fine-grained since they are too many from which to choose. As we saw in the last chapter, Kelly, on the other hand, is saying that any application of these demonstratives will fail anyway, since they cannot capture the experiential context of the referent – and this seems very much like Peacocke's claim that in conceiving of the demonstrative as targeting the reference we overshoot our mark.

experiences can have phenomenal content. In the next section I hope to show that there are good reasons for thinking that neither **(a)** nor **(b)** is the case, and that there is therefore no reason to accept **(3)**. We can, of course, deny **(3)** while maintaining **(1)** and **(2)**, which together are consistent with the position that some phenomenal states have conceptual content.

3.2 Critique

I will critique nonconceptualism along two lines. First, I will argue against **(a)** that certain aspects of our perception are best accounted for by an appeal to conceptual content: the phenomenological "shifts" which take place when perceiving ambiguous figures suggest that concepts enter into the content of visual states; and the phenomenological changes that take place when acquiring a language are best explained by allowing that concepts can enter into the content of auditory states. Second, I will suggest that there are good reasons for believing that thoughts and mental images have phenomenological aspects. But since thinking and imagining undoubtably involve conceptual mental states, it follows that there are some phenomenal states which aren't composed solely of nonconceptual content, and thus **(b)** is false.

A. Concept-dependent perception

In arguing for partial conceptualism it is not enough to establish that concepts "carve up" our experiences in order to make them memorable,

manageable, and intelligible. This much is patently obvious. I might think "This is a dog" or come to believe that there is a dog in front of me when I am in perceptual states with (nonconceptual) dog-like content, but the concept dog does not, on this rendering, enter into the content of my perception *itself*. It is, rather, part of the content of my *belief* about my perception. Such an understanding of the role of concepts maintains the split between the content of an experience and the beliefs which one comes to have concerning the experience. The function of concepts must be extended further so that, in addition to sorting, delimiting, and prioritizing the contents of phenomenal experiences, concepts actually form part of the perceptual content itself; it must be shown that perception is *imbued* with concepts. In this section I will present examples of experiences which I think meet this criterion for concept-dependence.

This part of my critique of nonconceptualism stems from what I think is the prima facie obviousness of the conceptualist thesis. By this I simply mean that it is difficult to doubt that our experiences of the world are structured at least in part by the concepts which we possess. Chris Peacocke makes the same point: "it is not clear that there is good reason for denying the overwhelmingly plausible view that we see things as trees or hear a sound as that of a car approaching."[146] When you stroll down a crowded sidewalk you have a vivid conscious perceptual experience of the scenery around you: you see tall buildings, colourful street

---

[146] Peacocke (1992) p. 123

displays, people, cars, and all the rest. There doesn't to be any sense in which what you *see* are *merely* ways of filling out space with bounded figures that afford you certain movements. Rather, your experiences are of objects and of events which are both structured and meaningful – *experienced* cars are more than just complex geometrical figures in your visual field. This conception of experience seems to me to be more than just a remnant of folk psychology.

One possible response a nonconceptualist would have ready at this early stage is that I am simply wrong about this: cars *just are* complex geometrical figures in our visual fields which we *judge* or *believe* to be cars, and while this judgement or thought plainly has a conceptual content it does not follow that the *experience* which the judgement is based upon is likewise conceptual. This is, after all, the basic contention of nonconceptualism: thought and experience are fundamentally different "activities" with different sorts of content. But there is certainly room for doubt here. It seems at least equally plausible that the ability to entertain the conceptual thought "There is a car over there", is made possible by an analogous conceptual experience *as of a car*, and that therefore this experience is concept-dependent.[147] Carruthers suggests that this latter picture of the relation

[147] Here I am endorsing a weakened version of what Noë (1999) calls the *dependency thesis*. For Noë, "perceptual experience is concept-dependent in the sense that when we have perceptual experience, we exercise our grasp of concepts." I differ from him in that I believe that while "as-of" perceptual experiences are concept-dependent it does not follow that *all* perceptual experiences are therefore concept-dependent.

between thought and experience is more intuitively plausible than the nonconceptualist view, which requires a sharp distinction between conceptual thought and nonconceptual experience: "phenomenologically, perception does not seem to be like that. I appear to *see* tables and chairs; not just filled spaces which I come to *believe* are tables and chairs."[148]

In other words, in seeing a car we do not seem to be making any inference to the best explanation upon an as-yet unlabelled perceptual experience. But this is the interpretation of experience that the nonconceptualist is recommending. As I will try to bring out below, I believe phenomenological evidence weighs against this explanation.

"As-of" experiences are a stronger case for conceptual content than Dretske takes them to be. I will consider two compelling cases in which this occurs: first, in instances of perceptual switching; second, in the acquisition of language. In both examples the phenomenological character of a purportedly nonconceptual perceptual state appears to be fundamentally changed by the *exercise* or the *acquisition* of concepts.

i. Perceptual switching – One promising line of argument for partial conceptualism attempts to establish that a perception "shared" by two people while they are both looking at the same scene might well turn out to be two dissimilar perceptions due to a difference in the concepts possessed by **A** and **B,** or

---

[148] Carruthers (2001) p. 130

that one person can have two distinct experiences of one scene depending on which concept he uses in "aiming" at it. This latter phenomenon is sometimes referred to as "perceptual switching." It is common to find instances in normal acts of perception where the possession of certain concepts fundamentally changes the content of a perception. But we shall have to do a little work in uncovering an uncontroversial example in support of this idea.

Perceptual switching occurs when a single image appears to represent one thing when considered under a certain concept and another thing entirely under another concept. A classic example is the rabbit/duck, which I might alternately consider as an instance of a duck, or a rabbit.[149] A nonconceptualist might deny that there is any phenomenological change in this picture whatsoever when the switch is performed. Though I can alternately see it is a duck or a rabbit, it is at least defensible that, *qualitatively,* the picture doesn't change. But on the other hand it is clear that *something* happens when I perform the perceptual switch. So what is it?

Carruthers rightfully suggests that it would be hard to explain the change as merely a change in our *beliefs*: "For *what* do I believe when I see the figure as a duck? Certainly not that it *is* a duck! But nor do I believe that it *is* a *picture* of a

[149] Note how this example is very different than the case we considered in the last chapter of seeing something alternately as an instance of a "square" or a "diamond." In that case, we could appeal to the "way" in which the shape was seen, which did not depend on our having concepts which named the shape but rather on the application of certain relations such as "is in symmetry with."

duck, either (for I know that it can equally be construed as a picture of a rabbit)."[150] Additional trouble for this response is that it is possible for me to see something *as a* rabbit without *believing* that I am seeing a rabbit, which would mean there is more to seeing something as a rabbit than having the occurrent belief "That's a rabbit", when undergoing a certain experience. This is evinced, as Carruthers suggests, by the fact that I can judge that there is a rabbit in front of me without *asserting* anything at all:

> **Whether the classification leads to the *belief* in the presence of rabbit will depend upon the rest of the person's cognitive state. But it still counts as an application of a concept, because the inferential connections to *mammal*, *eats carrots* and so on are all in place, ready to be accessed.**[151]

Further, as I have already mentioned above, it just seems obviously wrong that nothing about the images or sounds themselves change when considered under either of their aspects, as the 'belief' response would suggest. *Plenty* seems to change: the orientation of the animal, for instance, changes from left to right. Or, more simply: under one aspect it *looks to me* as if I am seeing a picture of a rabbit while under another aspect I see a picture of a duck. Carruthers therefore suggests that what changes when we consider the duck/rabbit image alternately as a duck and then as rabbit is the *concept* under which the image is organized: "We can, at will, make this switch its aspects back and forth – now a duck, now a

---

[150] Carruthers (2000), p. 131-132

[151] Ibid. p. 135-136

rabbit – depending on the concepts we deploy."[152] This seems to go through easier than the 'belief' response. I also believe he is right when he says that the difference which the conceptual reorganization brings about is a phenomenal difference. But let's turn to another example before spelling out the further implications of perceptual switching.

ii. Learning – Another phenomenon which supports a partial conceptualism is the effect that learning can have on the character of phenomenal perception. A commonly used example is that of hearing a foreign language. Consider two individuals, one, **A**, who has become fluent in the language, say Russian, while the other, **B**, has no knowledge of the language at all. The utterances of the Russian speaker will give rise to belief states and subsequent responses in the case of **A** that will not arise in the case of **B**. **A**, for instance, may experience the utterance as an exhortation, or perhaps as a query. Accordingly he will acquire certain beliefs about, say, the desires of the speaker which will afford him a range of possible behavioural responses. **B**'s auditory perception will not dispose him to form any such beliefs (barring an instance where he can "hear" the query in the quality of the speaker's voice, but even then the nature of the query is naturally not accessible to him). But if the experiences of **A** and **B** result in such vastly different belief states, then how can they be said to have the same perception at all? Of course most people would deny that they

[152] Ibid. p. 131

are having an identical experience; intuitively, A is hearing Russian and B is hearing gibberish. But then **A** and **B** must be having distinct experiences due to their possessing or lacking concepts; **A** has, loosely speaking, "Russian language" concepts which **B** does not possess.

But this example is not conclusive, since an obvious response is open to the nonconceptualist here. It could be argued that **A** and **B** are in fact hearing the same *pattern* of noises at the (nonconceptual) level in which auditory information is received, and it is rather only at the level of the *interpretation* of those noises – an activity which obviously involves concepts – where the content of the perceptions *seem* to differ. No doubt this retort is plausible at least in the case of the auditory perception of individual words. If the Russian speaker simply says "медведь," **A** will immediately recognize this as the Russian word for "bear"while **B** will have no idea what, if anything, was said. They will both, however, have heard the same sound, which suggests that, at this basic fine-grained, nonconceptual level of "raw sound," where the phenomenal character of the experience is putatively fixed, the content of their perceptions are identical. This experience will be nonconceptual in much the same way that the perception of a cloud contains shapes for which we have no concepts. We might say the cloud looks "fluffy," but this does not capture the fine-grained nature of the visual experience. Likewise, we might say that a word, though we might not even know its meaning, sounds "angry," but this concept does not exhaust the fine-grained

nature of auditory experience.

For reasons I will discuss shortly I do not think that this sort of response works; it does not follow that since they are each hearing the same sound they are having experiences with identical content. But nonetheless perhaps the case of hearing a single word is too contentious. If we can find an example that escapes this mire we would be better off for it. To that end, I believe instances of hearing *longer* strings of words in sustained speech can better support the conceptualist's thesis, since the point seems to go through without too much fuss about the nature of concept use. Consider: We can make a strong case that, while **A**'s "Russian language" concepts allow him to hear words where B hears only gibberish, these concepts also allow **A** to perceive the putatively nonconceptual auditory patterns in a fundamentally different way than does **B**. For when **A** hears the Russian speaker he perceives his speech to contain not only the words being spoken, but also minute pauses which occur between each word. The pauses, for **A**, mark the end of one word and the beginning of another, and their being perceived is critical for making sense of what is being said. But since **B** lacks the conceptual capacities to discern one word from another in a stream of Russian speech he will likely not perceive any pauses at all, since in spoken language there are seldom instances in normal speech of genuine silence, excluding deliberate pauses between sentences. **A** will instead likely perceive a steady and undifferentiated stream of sounds. This is no doubt why it often seems that those who speak

languages which we do not understand talk very rapidly. This suggests that possession of concepts can influence the phenomenal character of experience at the level of what nonconceptualists might call receptivity, and therefore experience must be at least partially conceptual.

In the case of language, the interpolation of pauses where there were none before seems to put the question beyond dispute: surely these are phenomenal differences if anything is! But a possible rebuttal comes from Michael Tye. He suggests that we can admit both that percepts can be organized via the concepts we possess and that different conceptual organizations result in different phenomenology while still insisting that the relevant concepts only exhibit a *causal* influence on the changes in the percept, changes that are themselves nonconceptual in nature. So in the case of the acquisition of language influencing the phenomenal character of sound, we can admit that "some aspects of phonological processing are sensitive to top-down feedback from the centers of comprehension."[153] "My claim," Tye says, "is that the phenomenally relevant representation of phonological features is nonconceptual, not that it is produced *exclusively* by what is in the acoustic signal."[154]

And likewise in the case of perceptual switching. We can of course allow that one cannot see something *as a rabbit* unless one possesses the concept rabbit

---

[153] Tye (2000) p. 61

[154] Ibid. p. 61

– and this is what conceptualists have rightly been insisting all along. But it does not follow from this that the representation which we are bringing under the concept rabbit is *itself* conceptual:

> **What happens in cases like these is that one has a sensory representation whose phenomenal content is then brought under the given concepts. Still, the concepts do not enter into the content of the sensory representation, and they are not themselves phenomenally relevant.**[155]

So the concepts "duck" and "rabbit" bring about a change in the phenomenology, but the change itself is a reorganization of the *non*conceptual content.

However tidy this explanation might be, I do not think it adequately deals with the challenges of perceptual switching and learning.[156] Recall that earlier, in arguing for partial nonconceptualism, I stated that perceiving something in a certain *way* is not dependent on possessing any concept. So for instance, I need

---

[155] Tye (1995) p. 140

[156] It is doubtful whether Tye himself endorses this argument anymore. In recent writings that Tye has warmed somewhat to the notion that experiences can contain conceptual content. As part of an E-symposium (http://host.uniroma3.it/progetti/kant/field/tyesymp.htm), in reply to Alex Byrne's paper "DON'T PANIC: Tye's Intentionalist Theory of Consciousness," Tye writes:

> **If I see a picture as a duck, then my visual state has a conceptual content, but it doesn't follow that it lacks any nonconceptual content. There are, it seems to me, many layers of perceptual content; and the possession by a perceptual state of one of these layers does not preclude it from having others.**

And in another recent paper entitled "On the nonconceptual content of experience," Tye remarks that his claim that experiences contain "robust" nonconceptual content is compatible with the further claim that some experiences also have conceptual content (p. 224, footnote), whereas in Ten Problems of Consciousness he argued that his view was incompatible with this claim.

not possess the concept "diamond" or "rabbit" in order to see something in that *way* in which a diamond or a rabbit is seen; the former would require only that I perceive certain symmetries rather than others which would typify a square. In the case of the latter, though it may be hard to imagine, I could just *happen* in perceiving certain relations between filled spaces in the perceptual array to see a figure in a way which could be described as "rabbit-shaped" though I lack the concept "rabbit." Alternately, following Dretske, we could call this an instance of being "thing-aware" of an object (seeing something which happens to be F) while lacking "fact-awareness" of it (seeing it as an F).

But I would like to suggest that seeing something merely in the *way* (in the technical sense in which Peacocke uses the word) in which a diamond or a rabbit is seen is not equivalent phenomenally to seeing a rabbit; an animal without conceptual abilities could see the first, but not the second. But since both of these experiences have the same nonconceptual content and are yet markedly different experiences they must therefore differ in their conceptual content. In other words, being in a state with the relevant nonconceptual content is necessary but not sufficient for 'seeing' these things *as* ducks or diamonds. Alva Noë similarly describes the experience of seeing geese flying overhead:

... I exercise my knowledge of (for example) what geese are, and what flying is. The experience is concept dependent because I could not have had just that experience of as of geese and flying if I did not have those concepts. This is not to say that one needs the concept of a goose to see a goose. The point is that one could not see a goose <u>as a goose</u> or <u>as flying</u> if

one lacked these concepts.[157]

Similarly, Peacocke suggests that if concepts did not enter into the content of our experiences "we would be unable to account for differences which manifestly exist." He goes on to give this example:

> **One such difference ... is that between the experience of a perceiver completely unfamiliar with Cyrillic script seeing a sentence in that script and the experience of one who understands a language written in that script. These two perceivers see the same shapes at the same positions . . . The experiences differ in that the second perceiver recognizes the symbols as of particular kinds, and sequences of the symbols as of particular semantic kinds.**[158]

Notice that in this example there is no change effected in the nonconceptual content (the "shape" and the "position") through possessing "Cyrillic script" concepts. It seems therefore that Michael Tye's claim that concepts merely cause a "reorganization" of the nonconceptual content is incomplete. I think the preceding shows that this is only part of the story, for even when no reorganization takes place, possessing concepts can effect a substantial *redescription* of the experience.

B. Conceptual content and 'inner' phenomenology

i. Mental imagery – Nothing seems clearer than that when we form a mental image there is an attendant phenomenology. Though they certainly lack the 'intensity' of, say, visual phenomenology, mental images nonetheless have

---

[157] Noë (1999), p. 257

[158] Peacocke (1992) p. 123

rich and varied qualities. And mental imagery is not confined to 'visual' images. I can also enjoy "auditory" images, as when I am trying to recall the melody of a song. There are even tactile, olfactory and gustatory mental images, though these have even weaker phenomenological components. Moreover, I can consult the qualities of my mental images in performing certain cognitive tasks. Mental rotation experiments are a classic example. The subject is presented with two complex geometrical objects and asked to "mentally rotate" them to discover whether or not they are two different views of the same object. Or suppose that you are asked to describe the shape which is enclosed within the capital letter 'A': "It seems entirely plausible that success in this task should require the generation of a visual image of that letter, from which the answer ('a triangle') can then be read off."[159] This suggests that mental images can best be described as internally generated phenomenal experiences.

I believe that issues surrounding mental imagery impact the debate with which we are concerned. For instance, Stephen Kosslyn argues that the natural comparison we make between mental imagery and perception is more than a simple analogy. Rather, imagery *is* quasi-perceptual. This is evinced by the fact that "most cortical areas used in visual perception are also used in imagery."[160] For instance, like perception, the generation of imagery involves activations of the

---

[159] Carruthers (1998b)

[160] Kosslyn et al. (2003) p. 109

occipital lobe, which contains numerous topographically mapped areas that support *depictive* or picture-like representations. A retinal image of a plus-sign will result in a plus-shaped pattern distributed across the occipital lobe. The occipital lobe constructs point-by-point isomorphic renderings of the image that falls upon the retina (in the case of perception) or of a stored memory representation (in the case of mental imagery).

Some evidence for the shared machinery of real world perception and mental imagery is that there are cognitive deficiencies which affect both in equal measure. In The Man Who Mistook His Wife for a Hat, Oliver Sacks describes the patient referenced in the book's title (Dr. P) as having problems with noticing things off to his left. Sacks wonders whether Dr. P might have the same trouble with mental images, and so asks him to describe the buildings that he would see while walking through a local square, first coming from the north and then from the south. In both cases, Dr. P failures to "notice" the buildings off to the left as he mentally walks through the square, even though the buildings he missed the second time around were the same buildings he correctly identified during his first walkthrough: "It was evident that his difficulties with leftness, his visual field deficits, were as much internal as external, bisecting his visual memory and imagination."[161]

The difference between perception and imagery is in the *functional* role of

[161] Sacks (1970) p. 15

these depictive representations that are, in part, generated by the occipital lobe. In perception, the representations constructed in the occipital lobe act as *input* to the two major visual pathways: One path, running down the inferior temporal lobe, is involved in object recognition. Visual memories are stored in this pathway, but in a non-topographic (we might say, descriptive or conceptual) form. The other path, which runs up to the posterior parietal lobe, is involved in specifying locations and orientations in space.[162] (We will return to these two visual pathways, often referred to respectively as the "what" and "where" paths, in the last chapter.) In generating mental imagery, on the other hand, the direction of causality runs opposite, so that the pattern on the occipital lobe is the output of activity in these two visual pathways. A mental image of a shape is rendered when a visual memory is activated "top-down," causing excitation in the occipital lobe, while images of spatial features of objects are created from spatial memory, causing excitation in the posterior lobes. But despite this difference between mental imagery and perception, once the images are created patterns in images can be processed much like the corresponding patterns of activation induced during perception.[163] The idea is something like this, then. In the case of forming a visual image of the letter 'A', a conceptual representation of the letter, stored in visual memory, is projected back through the visual system, thereby generating

---

[162] Ibid. p. 109

[163] Ibid. p. 109

activity in the occipital cortex, just as if a letter 'A' were being perceived. This activity is then processed by the visual system to yield a quasi-visual percept.

I take this to mean that (at least some) mental images are *conceptual* representations with phenomenal properties. Carruthers draws the same conclusion from an examination of Kosslyn's theory: "It is hard to see, then, how the quasi-percepts generated in visual imagery could fail to bear the imprint of concepts; and it does seem that concepts 'reach back down' through those levels of the visual system which are responsible for generating phenomenal consciousness."[164] This result certainly runs counter to the nonconceptualist thesis. Specifically, it causes problems for Dretske's claim that only representational systems with *systemic* indicator functions can give rise to phenomenal states, as well as Tye's claim that phenomenological states arise only at the level of sensory input to cognitive processes.

But I do not want to suggest that mental images have *no* nonconceptual content. On the contrary, I think it is likely that they do. As I have already mentioned, the evidence indicates that part of the mental image is constructed from the activities of the parietal lobe, which is responsible for coding spatial and other "agent-centred" properties of normal perceptual states. I have also suggested that those aspects of perceptual states which creatures exploit *qua* agents as they navigate the world are likely nonconceptual. I am therefore

[164] Carruthers (2000) p. 132

happy to grant that, insofar as mental images have spatial properties, etc. they have nonconceptual content. On the other hand, mental imagery seems to be most unlike perception in that the spatial properties of imagery are fragmentary and 'fuzzy'. Commenting on the evidence that patients with brain damage affecting their perception can sometimes retain normal capacities for forming mental images, Kosslyn writes that, while results like this clearly demonstrate that not all of the processes used in visual perception carry over to visual imagery, this is the sort of result we should expect from his theory:

> **... imagery relies on previously organized and stored information, whereas perception requires one to perform all aspects of figure-ground segregation, recognition and identification. We would not expect imagery to share "low-level" processes that are involved in organizing sensory input. In contrast, we would expect imagery to share most "high-level" processes, which involve the use of stored information.**[165]

This brings me to my second point. Much attention has been paid to the subpersonal, computational aspects of mental imagery. The phenomenology, I would argue, is less well understood. But despite this, no one denies that mental images have phenomenology, and that characterizing this phenomenology (and whether it is either misleading, accurate or unimportant to theories about the subpersonal aspects of mental imagery) is important to a complete theory of mental imagery.

My contention is, given that mental images have phenomenology, that

---

[165] Kosslyn et al. (2003) p. 320

phenomenology is, on the evidence of introspection, often characteristic of states which supposedly have only conceptual content, such as thought. For instance, I would argue that mental images at times exhibit a level *generality* which we do not find in perception, and which is characteristic of conceptual thought. The question as to whether there are such things as general mental images is quite an old topic in philosophy. Berkeley, for instance, was hostile to the idea. He thought that nothing could be more plain than that, when we imagine an object, it must have the same sorts of determinate properties that would be present in the perception of that object:

> **I can consider the hand, the eye, the nose, each by itself abstracted or separated from the rest of the body. But then whatever hand or eye I imagine, it must have some particular shape and colour. Likewise the idea of man that I frame to myself must be either of a white, or a black, or a tawny, a straight, or a crooked, a tall, or a low, or a middle-sized man. I cannot by any effort of thought conceive the abstract idea [of 'man'].**[166]

Berkeley is suggesting that it is impossible to form a mental image that would serve as an analogue to the concept "man." We can, of course, *think* about man abstractly, as a general category, but, try as we might, we can only form a mental *image* of some particular instance of a man. Thus, for Berkeley, inner and outer instances of 'seeing' are strongly analogous. Just as I cannot see an object out in the world without seeing it as having determinate characteristics, I also cannot 'see' a mental image without it being likewise determinate. For instance, nothing

[166] Berkeley (1710) p. 11

is *just* triangular. Whether I see a triangle or imagine one, the triangle that is before me must be isosceles, or equilateral, or scalene.

I do not think that Berkeley's argument is supported phenomenologically, for it seems to imply that all mental imagery is as clear, detailed and vivid as normal perception. But mental imagery is not as homogenous as Berkeley believes it to be; though there are certainly instances of mental imagery which strongly resemble perceptible objects, it seems to me that some of the mental images which I entertain could be described as fuzzy, inchoate, or indeterminate. I don't think that it is absurd to suggest that I can form a mental image of a zebra that is striped without it having any determinate number of stripes, or even that I can form a mental image of dog that is coloured without it being any determinate colour. In this way at least, mental imagery and perception are dissimilar. As H.H. Price writes in Thinking and Experience: "an image *can* have determinable characteristics without having any of the determinate characteristics falling under it."[167]

In his analysis of general images, Price goes on to say that the inchoate character of some mental images is connected with their evanescent or fleeting character. At times we might wish to "freeze" and scrutinize mental images, as when we are asked: "What is the shape of a Doberman's ears?" To answer this question we would probably do best to form a highly detailed percept-like image

---

[167] Price (1953) p. 287-288

of a Doberman's head facing in some direction. We might even rotate it to get a better look at its ears from different angles. But we rarely use mental images in this way. It is much more common to use mental images to *think with*, as Price says. In these cases it does not pay to scrutinize the image too closely: " ... we do not dwell upon them ... it would obstruct the process of thought if we did, or even arrest it altogether."[168] But one interesting effect of our not dwelling on mental images when we are thinking *with* them is that they aren't "nurtured" into becoming clear and distinct, percept-like pictures:

> **The full blooded, clear, detailed images which we sometimes have ... do not necessarily spring into existence ready made and complete. Sometimes they may. But often they *grow* more detailed and more clear 'before our mental eye.'**[169]

We should not be tempted to think that these fleeting images are determinate, that they only *seem* indeterminate because we do not attend to them. More accurately, the image itself is left incomplete: "the image does not have time to develop itself fully; it exists for so short a period that it is only half-formed, so to speak, by the time it vanishes."[170]

There is also a practical reason for why we do not attend to each mental image during thinking, since, often, an undeveloped image will serve as a better

---

[168] Ibid. p. 290

[169] Ibid. p. 290

[170] Ibid. p. 290

'general sign' for objects. For instance, if we wanted to utilize mental imagery when thinking about trees we would do well not to form a precise mental image of an apple tree since its specific characteristics will not resemble other tree species. Using this mental image might lead us to make errors in judgement. For instance, if you were to tell me that you owned a cabin in a densely wooded area without mentioning that it was located in Northern Ontario, I would be mistaken to imagine it as being surrounded by apple trees. So the more closely a mental image resembles a particular, the less closely it will resemble other particulars of the same class. If we accept that images should be fit to fulfil the role of standing for concepts, from this point of view "it would almost seem that a 'bad' image - schematic, sketchy, lacking in detail - is better than a 'good' one."[171]

On the other hand, Berkeley was right; we cannot form a precise mental image of a tree *in general*, since any well defined image, even a sort of Platonic line-drawing, will have some determinate characteristic of one sort or another. But there is a third option: Price suggests that a tree image that was on its way to becoming an apple tree, without having ever got there, would be much better suited for thinking imagistically about trees in general than either one of these.[172] To take another example, a hastily-formed image of a zebra in which no specific number of stripes coalesced would be a better stand in for the concept "zebra"

---

[171] Ibid. p. 275

[172] Ibid. p. 292

than a well formed mental image of a zebra with a determined amount of stripes. Price recognizes that the idea of general images may be repugnant to some philosophers. But this prejudice is based on a strong analogy between perception and imagery which doesn't square with the evidence of phenomenological evidence.

ii. Thought – Thought, as we have seen, is the paradigm of a state with conceptual content – in order to entertain a given thought I must exercise the concepts which are the constituents of that thought. It is also often assumed (less often argued for) that occurrent thoughts, on their own, are examples of purely intentional states, lacking in phenomenology. Both Dretske and Tye take this line, and so have little to say about the nature of conscious thought as it relates to phenomenology (presumably, both are committed to the view that thoughts have no phenomenology). Their silence on the subject is understandable; many theorists take the position that thoughts aren't *like* anything, in the relevant sense, to undergo. Thinking through math problems or sitting and idly brainstorming certainly don't *seem* to have any attendant phenomenology. I personally would find it difficult to pinpoint any clear-cut phenomenal shift that takes place when I engage in these thinking activities. So I am willing to admit that there is certainly an intuitive appeal to the position that thoughts aren't phenomenal, since it is difficult to indicate any phenomenological change which takes place when we go from thinking that p to thinking that q. By contrast, the phenomenological change

in the case of the rabbit/duck percept is plain to us.

However, others point out that, though some of the things we think do not have any particularly qualitative feel, other types of thoughts do have a phenomenal component. David Chalmers includes occurrent thoughts in his catalogue of phenomenally conscious experiences:

> **When I think of a lion, for instance, there seems to be a whiff of leonine quality to my phenomenology: what it is like to think of a lion is subtly different from what it is like to think of the Eiffel tower.**[173]

Although it may still be difficult to pinpoint just what the qualitative feel of certain thoughts are, it is nonetheless plausible that there is *something* it is like to undergo them. Like Chalmers, I think that something (perhaps barely detectable) does differ in our thought that *p* and our thought that *q*, something besides which parts of the world *p* and *q* are *about* (that is, something other than their intentional content); thought is not without its own, albeit subtle and strange, quality. And if this is true, it also stands to reason that the proper phenomenology of thought will be something other than the familiar types of perceptual phenomenology: gustatory, olfactory, visual, etc. But even if thought does not have its own *proper* phenomenology, the related claim made by Carruthers in his (1996) and (2005) that thought has phenomenology insofar as *inner speech* has a phenomenological component is much harder to deny. The argument goes like this: if inner speech is constitutive (and not merely expressive) of thought, and further if the auditory

---

[173] Chalmers (1996) p. 9-10

image of inner speech is conceptual (as I argued in the last section), it follows that thought *can* be phenomenally conscious; even though thought may not have proprietary phenomenology, it nonetheless acquires a *derivative* phenomenology when it takes the form of inner speech. Thus, even if we do not accept that thoughts have a propriety phenomenology (one which is entirely their own), it is still possible, as Peter Carruthers suggest, that they can at times acquire a derivative phenomenology in instances of inner speech, which are a sort of quasi-perceptual experience with thought content. I won't say any more about Carruthers' argument here.. But the falsity of nonconceptualism follows whether we accept that thought has proprietary or merely derivative phenomenology. Clearly, nonconceptualism entails that conceptual thinking *cannot* be phenomenally conscious. But if thoughts, which everyone agrees are paradigms of states with conceptual content, have any sort of phenomenological component, then it cannot be that phenomenology only belongs to nonconceptual perceptual states.

Let's briefly consider the case for a constitutive phenomenology of thought. In his article "The Phenomenology of Cognition Or What Is It Like to Think That P?", David Pitt presents an argument for a phenomenology of thought which examines certain abilities we have in relation to our thoughts, and concludes that we could only have these abilities if thoughts have an attendant phenomenology. Under normal conditions – barring any sort of cognitive

impairment – one has the ability, "consciously, introspectively and non-inferentially": to distinguish occurrent conscious thoughts from other occurrent conscious mental states (such a percepts); to distinguish between occurrent conscious thoughts; and to identify an occurrent thought as having the content it does. But, he argues, we could only have these abilities if each conscious thought had a phenomenal component which is of a different sort than that of perceptual phenomenology, which is "distinct from the phenomenology of other thoughts, and which is constitutive of its representational content."[174]

Pitt notes that these first two abilities are similar to what Dretske has called "non-epistemic" seeing which is equivalent to something looking$_p$. Recall that Dretske argues there is such a thing as simple seeing, as something looking$_p$ like a poodle since objects that cannot be visually identified may nonetheless be seen, that is, visually discriminated. Pitt is suggesting that when we introspect, we are likewise immediately acquainted with our thoughts in a way that is not unlike our immediate acquaintance with our perceptual states: "Simple perception is attentive experience of external objects; simple introspection is attentive experience of internal objects."[175] In either case, we can distinguish the thought or percept from other constituents of the mental environment purely on the basis of how it is experienced, that is, on the basis of its phenomenology alone, without

---

[174] Pitt (2004) p. 7

[175] Ibid. p. 10

necessarily having any beliefs about it.

That belief is not necessary is evinced by our ability to experience thoughts which we cannot identify: strange moods, fleeting thoughts, etc. Alvin Goldman suggests that Ray Jackendoff's example of "tip-of-the-tongue" thoughts demonstrates that "'conceptual' thought often occupies awareness or consciousness, even if it is phenomenologically 'thinner' than modality-specific experience.[176] When I try to think of something to say but cannot find the word, I am nonetheless phenomenologically aware of having a conceptual thought in mind. What I *cannot* find is simply the right phonological sound for the word I am seeking: "Entertaining the conceptual unit has a phenomenology, just not a sensory phenomenology."[177] Further, what it is like to entertain a given thought will elude description in conceptual terms in much the same way that the fineness of perceptual grain does. Descriptions of what it is like to have a thought will be outstripped by the richness of the thought's phenomenology. Pitt writes: "If one were to ask, 'Well, what is it like to think that the weather is changing?', I could only answer in the way I would if asked what it is like to see orange by someone who had never seen it."[178]

---

[176] Goldman (1993) p. 23

[177] Ibid. p. 23

[178] Pitt (2004) p. 31

2.3 Summation

These objections to conceptualism cumulatively make a case against **(3)**. There is no reason to believe that phenomenal states *just are* nonconceptual perceptual states poised to make an impact on the cognitive system. I have shown that the operations of the cognitive system can have a constitutive (and not merely causal) impact on the character of perceptual states. Further, the operations of the cognitive system *sans* perceptual input can create phenomenal states such as occurrent thoughts and mental images. This leaves us with **(1)** and **(2)** which, as I have mentioned are consistent with a mixed content theory of phenomenal mental states. Of course, it is possible to raise concerns about the coherence of the mixed content view – for how can the content of phenomenal perceptual states be both conceptual *and* nonconceptual in nature? In my final chapter I will attempt to answer this question, as well as provide some independent evidence that supports the view.

4. Mixed Content

"The visual field is that strange zone in which contradictory notions jostle each other ..."

- Maurice Merleau Ponty[179]

In the preceding two chapters I have tried to make a case for the content of phenomenal perception being neither wholly conceptual nor wholly nonconceptual in nature. In **Chapter 2**, I argued that a purely conceptualist view of perceptual content would not adequately deal with certain issues: the fineness of grain exhibited in perception, the interface between action and perception, and the perceptual lives of pre-linguistic infants and animals were my main points of contention. Using these examples, I attempted to show that in these situations creatures enjoy intentional mental states for which they need not have corresponding concepts. In **Chapter 3**, I followed the same programme in critiquing an exclusively nonconceptual view. I found objections to the thesis in the cases of perceptual switching and learning, which very strongly suggest that there is a conceptual strata in perceptual content. Attempts to account for these phenomena by invoking only nonconceptual contents were shown to be ineffective. Further, I argued that states such as thoughts and mental images

---

[179] Merleau-Ponty (1962) p. 6

which are commonly thought to be purely conceptual intentional states are actually tinged with phenomenology. If this is true then a theory which equates phenomenology with nonconceptual mental states must be mistaken. I hope to have made a convincing case (assuming we wish to maintain a representational theory of perception) for a "mixed content" theory of perceptual content, in which conceptual and nonconceptual content, while playing distinctive roles in the functional economy of mind, c*o-mingle* in phenomenal perception and thought.

4.1 Evidence from perceptual anomalies and lesion studies

A mixed content theory of phenomenal mental states presents us with an immediate problem, one which has likely been a factor in discrediting the theory for many: how can a single mental state contain *two different kinds* of content? But this problem only arises if we adopt a naive picture of, say, vision in which the visual scene is put together all at once by some sort of monolithic perceptual system. From the perspective of folk psychology this view seems self-evident, since when we open our eyes the world is simply *there* immediately. We do not have to wait as our perceptual systems divvy up the scene into regions of bounded space, define line and depth, bind color to shape, draws in different shades, and so on. Experience is simply *on* or *off.*

The truth is, of course, that perceptual systems are extremely complex and that they produce complex and multi-layered representations. Vision (but also

audition, touch, and the rest) "is exceedingly complex, too complex to operate all in one stage."[180] In the case of vision, representation "begins with information about light intensity and wavelength at the eye, and ends with a rich and many-layered representation of the visible scene," and between these, "processing occurs in a number of semi-independent modules."[181] The final scene presented to experience *is* cobbled together from the outputs of many perceptual subsystems operating at different stages, though we don't notice all of these strata coming together since they are laid down in an a unimaginably fast neuronal time-frame.

Some evidence for the stratified and semi-independent nature of perception is that, at times, a perceptual scene is generated that is internally inconsistent. M.C. Escher's *Waterfall* uses irregular perspectives to make it seem as if a waterfall is both above *and* below the plane of a water wheel which it is powering. Perceptual anomalies like this arise when different perceptual subsystems contribute inconsistent representations to the "pooled" representation. In this case, there seems to be a conflict between representations of depth and spatial relation; our eyes tell us both that the waterfall is further into the picture than the water wheel *and* that it empties onto the water wheel from above. These kinds of layers of perceptual representation – those that specify movement, depth, spatial relations, and so on – are likely nonconceptual, as Tim Crane

---

[180] Tye (2000) p. 70

[181] Ibid. p. 70

argues in "The waterfall illusion" (which concerns an altogether different illusion). Crane points out that if you stare at a waterfall or some other scene with dominant movement in one direction long enough and then turn your attention to some stationary scene, say, a wall, you will perceive movement in the opposite direction of the initial scene. And yet, at the same time, there is another sense in which the wall does *not* move at all, but appears to stay still. Crane believes that this aspect of the illusion poses a problem for the claim that the content of perceptual experience is composed of concepts. The illusion shows that we can have an explicitly contradictory experience, an experience of A and ~A. But we certainly cannot *believe* both A and ~A, since this would violate the Principle of Cognitive Significance outlined above. So it seems that these sorts of experiences cannot be belief-like.

I recognize that there are many instances of perception where conceptualization would simply get in the way of action. For this reason, nonconceptual content is often utilized for guiding action *here* and *now* in a way that requires minimal deliberation. Since perception involving motor responses must be extremely fast in order to keep up with a changing environment it would simply be a burden to weigh the information down with the sort of content that has rich inferential connections with thought. As Carruthers writes: "once the decision has been made to pick a particular berry ('That one is ripe') it does not much matter that it is a *berry* which one is picking – one just needs to guide the

fingers to frame the outline of a particular three dimensional object without crushing."[182]

For these reasons, I see no problem in accepting the result for which Crane is arguing, since it doesn't follow from this that *all* the strata of phenomenal states are nonconceptual, as both he and Tye would have it. As I have argued, I think there is a conceptual strata to phenomenology. In support of this thesis it should be noted that there are, in fact, other instances of perceptual anomalies in which the nature of the anomaly cuts *across* the divide between conceptual and nonconceptual phenomenal content (instead of being wholly contained in the nonconceptual content, as above). These disorders arise when one or the other of the parietal lobe visual stream – the "agent-centred" and action oriented pathway, and the temporal lobe stream – the "object-centred," conceptualizing pathway – is damaged while the other is still intact. Two such "mirror" perceptual pathologies strongly support a mixed content view of perceptual content.

Visual agnosia is often developed by people who have suffered brain damage localized in their temporal lobes. Visual agnosics are often unable to recognize or describe common objects, faces, pictures, or abstract designs despite their seemingly normal adequate perception of these things (visual agnosics can often see the object perfectly well, explore it in a normal way with their eyes, and

[182] Carruthers (2000) p. 159

even make attempts at describing it).[183] It is not simply that visual agnosics have difficulties finding words to describe what they are experiencing. It seems more accurate to say that they have lost the ability to "see as." D.F., a visual agnosic whose disorder is particularly well documented, developed visual agnosia following near-asphyxiation by carbon monoxide. As a result of her accident, D.F. has a number of perceptual deficits related to vision. For instance, D.F. was found to be unable to make a copy of a simple line drawing of an apple or an open book. The profundity of her impairment is clear when one looks at her copy of the apple picture; an undefined mess of horizontal and vertical lines which not even D.F. herself can make out. But when asked to draw those same objects from memory, D.F. performed as would a normal person.[184] This ability suggests that her stored scheme of "apple" is intact, and that her difficulty is instead in forming a conscious perception of what she is looking at *as an apple*, rather than a difficulty in matching her perception to memories.

Another curious facet of certain visual agnosics is that when they mistakenly identify objects they often confuse them with other things that are semantically or conceptually linked to that object. For instance, when shown a knife, a patient reports "It's a plate!" Other agnosics might be able to identify the *category* to which an item belongs. Neurologist Jason W. Brown relates the case

---

[183] Brown (1988) p. 184

[184] Carruthers (2000) p. 161

of a mathematics professor with agnosia who "identified a slide rule as 'something for measuring, for calculation'" and "a Star of David as 'something holy.'"[185] These abilities, says Brown, "indicate that *early perception involves the selection of a target through a representational network or a memory organized about relations of conceptual proximity.*"[186]

But despite these difficulties, it is striking that visual agnosics can interact with the object as they normally would, even though they claim not to know what the object is. For instance, in reaching for a pen they will adopt the handgrip which they commonly use for holding a writing utensil, despite that they are unaware of what it is they are grasping. And when asked to identify the orientation of a line, visual agnosics do poorly, although they perform almost normally when asked to post a letter through a similarly shaped slot oriented at random angles.[187] I would venture that these traits of visual agnosics are best explained as resulting from a deterioration of the conceptual content of their visual experiences.

The mirror syndrome, in which the temporal lobe is damaged while the parietal lobe is left intact, is called optic ataxia, also sometimes referred to as *Bálint-Holmes syndrome*. Subjects who suffer from this perceptual pathology

---

[185] Brown (1988) p. 185

[186] Ibid. p. 184

[187] Carruthers (2000) p. 161

have no difficulties in identifying or reporting on the orientation of the objects in their visual field. However, they are at pains to interact with these objects properly: "they can be very poor at moving their hand in the direction of a target; they often cannot orient their hand correctly to pass it through a variably-oriented letter box slot; and they are poor at reaching out and grasping objects."[188] Once again, it is hard to make sense of this disorder except to say that the conceptual content of their experience remains intact while, at the level of nonconceptual content experience is indeterminate and, therefore, deceptive. Christopher Peacocke has a similar explanation for the disorder:

> **. . . the representations which control the limbs do use scenario content [a type of nonconceptual content], but . . . in the case of optic ataxia the contents of visual experience are *inaccessible* to the motor control systems.**[189]

The mixed content view of phenomenal states provides the most coherent explanation of these twin disorders, and seems to be more consistent with the dual-function theory of vision than either conceptualism or nonconceptualism.

4.2 Philosophical justifications

Though I have argued that the content of experience is partly conceptual and partly non-conceptual, I do not intend for this to mean that these two types of

---

[188] Ibid. p. 161

[189] Peacocke (1992) p. 128

content have radically different metaphysical origins. But, if I am right, the initial debate has hinged on concepts and percepts having disparate metaphysics: the first, "mental," the second, "worldly." Having arrived at the mixed content view, we must ask ourselves with hindsight if the initial terms of the debate made any sense at all. I am more sympathetic with Uriah Kriegel, who says that the conceptual/nonconceptual divide simply highlights

> **. . . that the functional role your experience plays in your mental life involves the mobilization of a mixture of capacities, some of which are recognitional [conceptual] and some merely discriminatory [nonconceptual].**[190]

Of course, our ability to *recognize* is closely tied with our conceptual repertoire, while our ability to *discriminate* is more fine-grained than this. But this does not mean that one or the other must alone establish phenomenal character.

If there *is* a salient difference between these two functions of our experience, it is, as Adrian Cussins suggests, the governing norms for either sort of content. I end with a brief discussion of his ideas, because I think they represent a promising way of moving forward in a positive account of a dual content theory.

Cussins' strategy in developing a theory of content differs from most others, since he does not proceed from the assumption that all content must present the world in such a way that it can be judged true or false. Instead, he

---

[190] Kriegel (2004) p. 12

begins by examining a kind of normativity and asking whether there is a type of content which it governs. If there is a kind of content that is guided by those norms "then we can explain how the world is presented in this content in terms of the structure that is necessary for explaining the norm-governed relations between these norm-governed contents."[191]

The guiding norm of conceptual experiences – which Cussins categorises, alongside judgements and beliefs, as *thoughts* – is truth. This is because when we conceive of the world through thought we conceive it "in terms of the structure which is necessary to characterize the truth of the thought or its truth-governed relations – inferential relations – to other thoughts."[192] Conceptual content, therefore, is that way in which the world is conceived that can render thoughts true or false. This is another way of saying that conceptual experiences present the world in a way that can be used to justify beliefs. And just as beliefs and judgements are constructed from discrete and recombinable components, so too are conceptual experiences; in a conceptualized world we are presented with the world as a realm of reference, a "realm of objects, properties, states of affairs, . . ., with respect to which the *truth* of thoughts is determined.[193]

As I have argued, there are also nonconceptual contents in experiences. If

---

[191] Cussins (1990) p. 153

[192] Ibid. p. 152

[193] Ibid. p. 153

we are to motivate this different sort of content, says Cussins, then it will have to be a different kind of content in which the world is not presented as a truth-maker (in other words, we are looking for a presentation of the world that isn't *thought-like*). Being representational, these contentful states must still be capable of going awry or *mis*representing. But the standards for misrepresentation ought to be something *other* than truth and falsity.

The form of guidance that governs nonconceptual content is so familiar, Cussins says, that we rarely ever pay attention to it. Imagine entering a room and trying to find one's way from one side to the other. When navigating the room we are not principally guided by the norm of truth. I do not, for instance, ask myself whether it is true that there is a chair in my way. Rather, I am guided by a sense of my body in relation to the solid objects around me, by the "activity-space" around me.[194] As I bump up against furniture I reorient myself, correct my path, etc.:

---

[194] It is interesting to note that Cussins here implicitly acknowledges an affinity that his distinction has with a line of thought from Heidegger's Being and Time. He remarks:

"Just as – when things are functioning well – the tool is not given to the builder as an *object of thought*, so the structured space of the room is not typically an object for subjects in the room (it is not part of the subject's realm of reference.)" [p. 155]

This echoes Heidegger's remarks about the distinction between two different attitudes which we can take to things in our environment. An object of thought roughly corresponds to what Heidegger called the present-at-hand. Making something "present" is something like conferring upon it the status of an object, or, we might say, it is a way of bringing it into the realm of reference. Conversely, when something is ready-to-hand it is not given as an object of thought to be analyzed, but rather as a thing of interest to a project in which the individual is engaged. I simply offer this observation as an instance of the rich connections

"These 'bumpings' are forms of guidance, and so they manifest candidate norms, not with the status of the norm of truth but fully normative nonetheless."[195] That is, the things in the room are cognitively salient insofar as they call to us, guiding us either toward or away from them.

The structure of the room, says Cussins, dictates the structure of my activity in a way that could remain identical despite my having one or another intention in proceeding through the room. Recalling J.J. Gibson, Cussins calls this the structure "of afforded paths or *trails* through the environment of the room" which fix "a distinction between skilled and unskilled, or competent and incompetent activity in the environment."[196] My judgements are evaluated with respect to truth, while my actions reflect my level of skill in mediating the environment. Cussins calls this intention-free normativity "mundane normativity".

In contrast to the world presented as a *realm of reference*, nonconceptual content presents the world as a *realm of mediation.* In the room, I might be presented with a realm of reference that includes people and furniture, as well as perceived relations between people both spatial and social, etc. The realm of

---

between this area of analytic philosophy and the continental tradition, which, for reasons of space, aren't pursued here.

[195] Ibid. p. 154

[196] Ibid p. 154-155

mediation, on the other hand, will consist "of the trails that distinguish patterns of afforded activity from patterns of resisted activity, which guide me – as a skilled traverser of rooms – as I cross to the other side."[197] As evinced by the example of optic ataxia – which we might call a disorder of nonconceptual content – the role of nonconceptual content is to assist us in skillfully navigating the environment, in making accurate discriminations and acting accordingly on them.

I think there is something to this. It is certainly true that there are two forms of guidance: "the kind of guidance that is provided by propositional judgement, which is employed, for example, in practical and theoretical reasoning, and more mundane kinds of guidance employed in everyday getting-about."[198] It is also true that these activities are guided by different norms. We might talk about a *skillful* use of reasoning in a legitimate way, but the *governing* norm of reason remains truth. And it seems odd to say that we can move about the world in a way that is true or false, though we certainly can *think* about it in ways that are true or false. It seems plausible therefore that thoughts and actions present the world in fundamentally different ways, with different sorts of contentful states governed by different norms.[199]

---

[197] Ibid. p. 155

[198] Ibid. p. 158

[199] It has been suggested that accounts of nonconceptual content rely on two different notions of nonconceptual content. One "sort" of content is the type that we get from fineness of grain arguments. This type of content is non-doxastic

Do the infant and the adult see the same thing when looking at the world? At the nonconceptual level of what Dretske has called "thing-awareness," their experiences might well be identical. Perhaps they are looking at the same area of filled space, perceiving the same symmetries, etc. – we could say, along with Peacocke, that they have the same "scenario content" and "proto-propositional content." Or, with Cussins, we could say that the trails of activity that are afforded to them by the perception of the objects in their environment are identical, though their skills and intentions for dealing with that environment might differ. But at the same time that the visual systems of both the adult and the infant are constructing these low-level representations, these basic perceptual discriminations and detecting of affordances, their visual systems are "*also* trying to impose concepts on the representations being constructed in attended-to-regions

---

– that is, it is richer and more fine-grained than the doxastic component of perception is able to capture – and yet it is representational. The other "sort" is the one that is used in the guidance of behaviour, and is governed by mundane norms of skillful activity. But note, however, that the "fineness of grain" comes out of a *critique* of conceptualism, whereas the "behaviourist" notion of nonconceptual content is an attempt at a *positive* account of nonconceptual content. The fineness of grain – presented, for instance, in Peacocke (2001) argument simply tells us that there *are* ways in which we represent the world that do not necessarily involve concepts. On the other hand, the "behaviorist" account of nonconceptual content simply answers these questions: "How do these nonconceptual contents represent the world, and, granted that they aren't doxastic and aren't used in forming beliefs about the world, what sort of tasks *are* they useful for?" Plausible answers for these questions are found in Cussins (1990), Hurley (1998), and Bermudez (1998).

of the visual field, on a 'best-fit' basis."[200]

[200] Carruthers (2000) p. 136

Bibliography

Articles

Bermudez, J. "Nonconceptual self-consciousness and cognitive science," *Synthese* 129: 129-149, 2001.

Block, N. "On a confusion about a function of consciousness," *The Behavioural and Brain Sciences,* 18 (1995) 227-247

Brewer, B. "Foundations of perceptual knowledge," *American Philosophical Quarterly*, vol:34 iss:1 pg:41. 1997.

Carruthers, P. "Consciousness and concepts," *Aristotelian Society*: Supplementary Volume, no. 66 (suppl.), pp. 41-59, 1992

Carruthers, P. "Conscious thinking: language or elimination?" *Mind and language* 13(4). 1998a.

Carruthers, P. "On Fodor's problem," *Mind and Language*, 18(5), pp. 502-523. 2003.

Carruthers, P. "Thinking in language?: evolution and a modularist possibility," in Carruthers, Peter and Boucher, Jill, Eds. Language and thought: interdisciplinary themes, chapter 5, pages pp. 94-119. 1998b.

Carruthers, P. "Who is blind to blindsight?" *Psyche*, 7(04), March 2001

Clark, A. "Connectionism and cognitive flexibility," in Essays on Nonconceptual Content. Gunther Y. (ed.) MIT Press: Cambridge. 2003.

Clark, A. "Visual experiences and motor action: Are the bonds too tight?" *The Philosophical Review*, Vol. 110, No. 4 (Oct., 2001), 495-519

Crane, T. "The waterfall illusion," *Analysis* vol. 48 (1988) 142-147

Cussins, A. "Content, conceptual content, and nonconceptual content," in *The Philosophy of Artificial Intelligence*, M. Boden (ed.) Oxford University Press: Oxford. 1990. p. 133-163.

Davidson, D. "A coherence theory of truth and knowledge," reprinted in Ernest LePore, ed., *Truth and Interpretation: Perspectives on the Philosophy of Donald*

Davidson. Basil Blackwell: Oxford. 1986. p. 307-319.

Dennett, D. "The path not taken," *Behavioural and Brain Sciences*, 18 (1995) 252-53.

Dennett, D and Kinsbourne, M. "Time and the Observer" *Behavioural and Brain Sciences*, 15 (1992) 183-247.

Goldman, A. "The psychology of folk psychology" *Behavioural and Brain Sciences* 16: 15-28 (1993).

Goodale, M.A. and Humphrey, G.K. "Separate visual systems for action and perception," in Blackwell Handbook of Perception, Goldstein, E.B. (ed.) 311-343. 2001.

Goodale, M.A and Milner, A.D. "Separate visual pathways for perception and action" *Trends in Neurosciences*, 15. Reprinted in Essential Sources in the Scientific Study of Consciousness, Baars, B., Banks. W. and Newman, J, (eds.) 175-185. 2003.

Hamlyn, D. "Perception, sensation and non-conceptual content," *The Philosophical Quarterly,* Vol. 44, No. 175 (Apr., 1994), 139-153

Hanson, N. "From patterns of discovery," in Perception, R. Schwartz, ed. Blackwell Publishing: Malden. 2004. p. 292- 305

Hanna, R. "Kant and nonconceptual content," *European Journal of Philosophy,* 13 (2005): 247-290.

Hurley, S. "Overintellecualizing the mind," *Philosophy and Phenomenological Research*, Volume 63. Number 2, September 2001. p. 423-431

Kosslyn, S., Ganis, G. and Thompson, W. "Mental imagery: Against the nihilistic hypothesis," *Trends in Cognitive Sciences* vol:7 iss:3 pg:109 -111. 2003.

Kosslyn, S., Thompson, W. and Alpert, M. "Neural systems shared by visual imagery and visual perception: A positron emission tomography study," *Neuroimage,* 6, 320-334. 1997.

Kriegel, U. "Perceptual experience, conscious content, and non-conceptual content," *Essays in Philosophy*, Vol. 5 No. 1, January 2004

Kriegel, U. "The new mysterianism and the thesis of cognitive closure," *Acta Analytica: Philosophy and Psychology*, vol. 18, no. 30-31, pp. 177-191, 2003.

Jackson, F. "Epiphenomenal qualia," *Philosophical Quarterly*, 32 (1982), 127-36.

Martin, M. "Perceptions, concepts, and memory," *Philosophical Review*, vol. 101, no. 4 (1992) p. 745-764.

McDowell, J. "Precis of Mind and World," *Philosophy and Phenomenological Research.* Vol. LVIII, No. 2, June 1998.

McDowell, J. "The content of perceptual experience" *The Philosophical Quarterly*, Vol. 44, No. 175 (Apr., 1994), 190-205

Nagel, T. "What is it like to be a bat?" *The Philosophical Review* LXXXIII, 4 (October 1974): 435-50.

Noë, A. "Is perspectival self-consciousness non-conceptual?" *The Philosophical Quarterly*, Vol. 52, No. 207 (Apr., 2002), 185-194

Noë, A. "Thought and experience," *American Philosophical Quarterly*, Vol. 36. No. 3 (July 1999) 257-265

Peacocke, C. "Does perception have a nonconceptual content?" Journal of Philosophy 98, 239-264. 2001.

Peacocke, C. "Nonconceptual content defended," *Philosophy and Phenomenological Research*, Vol. LVIII, No. 2, June 1998. p. 381-388

Peacocke, C. "Scenarios, concepts, and perception," in Essays on Nonconceptual Content. Gunther, Y. (ed.) MIT Press: Cambridge. 2003.

Pitt, D. "The phenomenology of cognition or what is it like to think that P?" *Philosophy and Phenomenological Research*, Volume 69, Number 1, July 2004.

Pylyshyn, Z. What the mind's eye tells the mind's brain: A critique of mental imagery
Psychological Bulletin, 80, 1-24. 1973.

Roy, J-M. "Phenomenological claims and the myth of the given," in The Problem of Consciousness. Thompson, E. (ed.) University of Calgary Press: Calgary. 2003.

Taylor, C. "The validity of transcendental arguments," *Proceedings of the Aristotelian Society*. Vol. 79.

Thompson. E. "Look again: Phenomenology and mental imagery," *Phenomenology and the Cognitive Sciences*. Special issue on phenomenology and heterophenomenology, edited by Alva Noe. Penultimate draft.

Von Uexkull, J. "A stroll through the worlds of animals and men," in Instinctive Behaviour, ed. K Lashley. International Universities Press. 1934.

Weiskrantz, L., Warrington, E. K., Saunders, M. D., and Marshall, J. "Visual capacity in the hemianopic field following a restricted occipital ablation," *Brain*, vol. 47 (1974)

Weiskrantz, L. "Some contributions of neuropsychology of vision and memory to the problem of consciousness," in *Consciousness in Contemporary Science*, Marcel A., and Bisiach, E. (eds.) Oxford University Press: New York. (1988)

Zahavi, D. "First-person thoughts and embodied self-awareness: Some reflections on the relation between recent analytical philosophy and phenomenology," *Phenomenology and the Cognitive Sciences* 1: 7-26, 2002.

Books

Armstrong, D. A Materialist Theory of Mind. Routledge & Kegan Paul. New York. 1968.

Bennett, M.R. and Hacker, P.M.S. Philosophical Foundations of Neuroscience. Blackwell Publishing: Malden. 2003.

Berkeley, G. Principles of Human Knowledge. Oxford University Press: Oxford. 1710.

Bermudez, J.L. The Paradox of Self-Consciousness. MIT Press: Cambridge. 1998

Brewer, B. Perception and Reason. Oxford University Press: Oxford. 1999.

Brown, J.L. Life of the Mind: Selected Papers. Lawrence Erlbaum Associates: Hillsdale. 1988.

Carruthers, P. Phenomenal Consciousness: A Naturalistic Theory. Cambridge

University Press: Cambridge. 2000.

Carruthers, P. Consciousness: Essays from a Higher-Order Perspective. Oxford University Press: Oxford. 2005

Chalmers, D. The Conscious Mind. Oxford University Press: New York. 1996.

Clark, A. Being There. The MIT Press: Cambridge. 1997.

Dennett, D. Brainstorms: Philosophical Essays on Mind and Psychology. Bradford Books. 1978.

Dennett, D. Consciousness Explained. Little, Brown, and Company: Boston. 1991.

Dennett, D. Sweet Dreams: Philosophical Obstacles to a Science of Consciousness. The MIT Press: Cambridge. 2005.

Descartes, R. Meditations on First Philosophy: With Selections from the Objections and Replies. John Cunningham. (eds.) Cambridge University Press: Cambridge. 1996.

Dretske, F. Knowledge and the Flow of Information. The MIT Press: Cambridge. 1981.

Evans, G. The Varieties of Reference. Oxford University Press: Oxford. 1982.

Fodor, J. Representations. The MIT Press: Cambridge. 1981.

Gibson, J. The Ecological Approach to Visual Perception. Lawrence Erlbaum Associates: Hillsdale. 1979.

Gunther, Y. (ed.) Essays on Nonconceptual Content. MIT Press: Cambridge. 2003.

Hamlyn, D.W. In and Out of the Black Box. Basil Blackwell Ltd: Oxford. 1990.

Heidegger, M. Being and Time. Harper & Row: San Francisco. 1926.

Hurley, S. Consciousness in Action. Harvard University Press. Harvard University Press: Cambridge. 1998

Jackendoff, R. Consciousness and the Computational Mind. The MIT Press: Cambridge. 1987

Kelly, S. The Relevance of Phenomenology to the Philosophy of Language and the Mind. Garland Publishing: New York. 2001.

Kant, I. Critique of Pure Reason. Guyer, P. and Wood, A. (trans.) Cambridge University Press: Cambridge. 1998.

Kosslyn, S. Image and Brain. The MIT Press: Cambridge. 1994.

McDowell, M. Mind and World. Harvard University Press: Cambridge. 1994.

McGinn, C. The Problem of Consciousness. Blackwell Publishing: Oxford. 1990.

Merleau-Ponty, M. Phenomenology of Perception. Routledge: London. 1962.

Nagel, T. The View from Nowhere. Oxford University Press: Oxford. 1986.

Nagel. T. Mortal Questions. Cambridge University Press: Cambridge. 1979.

Price, H. Thinking and Experience. Hutchinson's University Library: New York. 1953.

Quine, W. From a Logical Point of View. Harvard University Press: Cambridge. 1953.

Reid, T. Essays on the Intellectual Powers of Man: A Critical Edition. (ed. Derek R. Brookes) Edinburgh University Press: Edinburgh. 2002.

Sellars. W. Science, Perception and Reality. Routledge and Kegan Paul: London. 1963

Strawson, P. The Bounds of Sense. Methuen & Co. Ltd: London. 1966.

Tye, M. Consciousness, Color, and Content. The MIT Press: Cambridge. 2000.

Tye. M. Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind. The MIT Press: Cambridge. 1995.

Web Resources

Tye, M. "To PANIC or not to PANIC? - Reply to Byrne" 2002. Retrieved September 20, 2006, from http://host.uniroma3.it/progetti/kant/field/tyesymp_replytobyrne.htm