

Experience, Intentionality, and Concepts

Essays in the Philosophy of Perception

Anders Nes

Dissertation submitted in partial fulfilment

of the requirements for the degree Ph.d.,

Faculty of Humanities,

University of Oslo,

2007

© Anders Nes, 2008

*Series of dissertations submitted to the
Faculty of Humanities, University of Oslo
No. 336*

ISSN 0806-3222

All rights reserved. No part of this publication may be reproduced or transmitted, in any form or by any means, without permission.

Cover: Inger Sandved Anfinsen.
Printed in Norway: AiT e-dit AS, Oslo, 2008.

Produced in co-operation with Unipub AS.
The thesis is produced by Unipub AS merely in connection with the thesis defence. Kindly direct all inquiries regarding the thesis to the copyright holder or the unit which grants the doctorate.

*Unipub AS is owned by
The University Foundation for Student Life (SiO)*

Acknowledgements

I am grateful to my advisor, Carsten Hansen, for valuable feedback on and support for the project. Carsten also provided its original inspiration, as my interest in the issues it addresses first was stirred at one of his eye-opening seminars during my undergraduate days. I would also like to record a long-standing debt to Olav Gjelsvik for many helpful conversations over the years at the department at Oslo.

I am grateful to the Faculty of Humanities at the University of Oslo for electing me to a research fellowship, and to the Department of Philosophy, Classics, History or Art and Ideas for providing me with an excellent environment to work in. Special thanks, on this score, to the terrific crew at NT325. The Institute of Philosophy of the School of Advanced Study, London, gave me the opportunity to come as a visiting fellow for the year 2006-7. I thank the Institute for a rewarding visit, and the Faculty of Humanities at Oslo for helping to fund it.

I have profited from conversations with or comments from a great number of colleagues including, though not limited to, Keith Allen, Wylie Breckenridge, Bill Brewer, John Campbell, Herman Cappelen, Tim Crane, Colin Johnson, Øystein Linnebo, Richard Price, Susanna Schellenberg, Ralph Wedgwood, and Tim Williamson. Except for essay 6, earlier versions of the papers in the dissertation have been presented at graduate conferences at Columbia-NYU, Edinburgh, Harvard-MIT, and Nottingham and at

seminars at London's Institute of Philosophy, Oslo, and Oxford. I am grateful to my audiences for valuable feedback. Thanks in particular to Paul Katsafanas, Geoff Lee, and Matthew Nudds who were my formal commentators at Harvard-MIT, Columbia-NYU and Edinburgh, respectively.

Cordial thanks, finally, to family, and friends, and Shirin, in particular, for their patience, love and support during my work on this project.

Preface

This dissertation collects six recent essays of mine in the philosophy of perception, some of which branch over into general philosophy of mind. Some of the essays (numbers 2 and 3 in particular) deal with issues that are closely related to themes dealt with in my Oxford University thesis, *Content in Thought and Perception*.¹ However, they in each case give new arguments for the conclusions they reach, and stand on their own feet. I discuss the relationship between these essays collected here and *Content in Thought and Perception* in section 7 of the Introduction.

¹ Nes (2006).

Contents

<i>Acknowledgements</i>	iii
<i>Preface</i>	v
Introduction	1
1 <i>The notions of perception and perceptual experience</i>	2
2 <i>Essay 1: Weak representationalism</i>	8
3 <i>Essays 2, 3, and 4: Conceptualism and nonconceptualism</i>	12
3.1 <i>Conceptualised states</i>	16
3.2 <i>Essay 2: Conceptualism, recognitional concepts and indiscriminability</i>	24
3.3 <i>Essay 3: Conceptualism, self-conscious thought and systematicity</i>	26
3.4 <i>Essay 4: Content conceptualism and diaphanousness</i>	29
4 <i>Essay 5: Phenomenal concepts in introspection and perception</i>	30
5 <i>Essay 6: Intentionality and the mark of the mental</i>	32
6 <i>Interconnections</i>	34
6.1 <i>The bearing of weak representationalism</i>	34
6.2 <i>The bearing of state nonconceptualism</i>	37
6.3 <i>The bearing of the incompatibility of content conceptualism and (D)</i>	42
7 <i>Relation of this work to Content in Thought and Perception</i>	48
7.1 <i>Précis of Content</i>	49
7.2 <i>Relations</i>	58
8 <i>Concluding remark on contribution</i>	62
1 Looks and Representational Content: A Defence of the Orthodoxy	64
1 <i>The commitments of representationalism</i>	67
2 <i>Misleading but not misrepresenting</i>	71
3 <i>Looks-indexing I: Looking a certain way</i>	75
4 <i>Are some qualitative looks statements to be understood non-comparatively?</i>	87
5 <i>Looks-indexing II: Demonstrable looks</i>	93
5.1 <i>Contextual paraphrase in terms of looking a certain way</i>	99

5.2	<i>Demonstrable looks as sense-data</i>	99
5.3	<i>Demonstrable looks as physical particulars</i>	102
5.4	<i>Demonstrable looks as universals</i>	108
6	<i>Seeing and representational content</i>	116
7	<i>Conclusion</i>	121
2	Conceptualism, Recognitional Concepts, and Indiscriminability	123
0	<i>Introduction</i>	123
1	<i>McDowell and Brewer on phenomenal sorites series</i>	129
2	<i>The First Rule: A difficulty</i>	132
3	<i>The Goodmanian Rule and the sorites sceptical argument</i>	137
3	Conceptualism, Self-Conscious Thought, and Systematicity	144
0	<i>Introduction</i>	144
1	<i>The objection from fineness of grain, and some responses</i>	148
2	<i>Concepts and self-conscious thought</i>	154
3	<i>Systematicity and its failures in perceptual experience</i>	166
4	<i>Objections</i>	178
5	<i>Conclusion</i>	188
4	A Problem about the Compatibility of Content Conceptualism with a Certain Thesis of Diaphanousness	190
0	<i>Introduction</i>	190
1	<i>The problem</i>	195
2	<i>Two objections to (2), and replies</i>	200
3	<i>Options for the content conceptualist</i>	202
5	Object-Directed Knowledge Arguments and Phenomenal Concepts in Perception	211
0	<i>Introduction</i>	211
1	<i>Jackson's knowledge argument and the perspectivalist response</i>	213
2	<i>An experience-directed argument for phenomenal concepts in introspection</i>	217
3	<i>An object-directed argument for phenomenal concepts in perception</i>	223
4	<i>Does Mary already know what tomatoes are like, colour-wise?</i>	229

5	<i>Is the new belief O Mary acquires interestingly different from the experience-directed belief E?</i>	232
6	<i>Conclusion</i>	239
6	Are Only Mental States Intentional?	241
0	<i>Introduction</i>	241
1	<i>A problem</i>	242
2	<i>A dismissive response</i>	243
3	<i>Failures of substitution and existential presupposition</i>	246
4	<i>Directedness upon impossibilities</i>	247
5	<i>Varieties of substitution failure</i>	249
6	<i>Strong and stronger substitution failures</i>	251
7	<i>Conclusion</i>	254
	<i>Bibliography</i>	256

Introduction

For almost as long as there have been philosophers, there have been philosophers interested in perception. Perhaps this is not surprising. Perception is the interface between the world and us. It informs our thought about and knowledge of the world, and guides our action in it. To understand our place in the scheme of things, perception is one of the things we need to understand.

The dissertation collects six essays that fall, more or less squarely, in the philosophy of perception. The first four discuss the nature of perceptual experience, with special reference to the notions of representational content and conceptualisation. The fifth compares perceptual and introspective judgement. The sixth addresses an issue in general philosophy of mind – viz., whether intentionality is the mark of the mental – but does so in a way that is found to connect with issues about sensation and perception.

If there is a unifying question running through the essays, it is what the similarities and differences may be between perceptual experience and thought. I argue that perceptual experiences are similar to thoughts in that they both are states with intentional or (as one might also say) representational content. At any rate, I defend this claim against some recent interesting objections. I argue that perceptual experiences are dissimilar to thoughts in that they are not conceptualised states of mind. At least, I argue that the

contrary view faces significant problems, hitherto unacknowledged by its proponents. The connections between experience and intentionality, and between experience and concepts, are thus the two major themes of the dissertation.

There are also various supplementary themes. One is the connection between the putative diaphanousness or transparency of experience and its representational content. Another the role of so-called phenomenal concepts in introspective and perceptual judgement. A third theme hovers in the background through much of the dissertation, and is explicitly addressed towards the end. It concerns the highly general question about the role of the notion of intentionality in an account of the mind.

In this Introduction, I first explain the notions of perception and perceptual experience as deployed here. I go on to give an overview of the themes and arguments of the six essays to follow. Special attention is devoted here to the contrast (or rather contrasts) between conceptualism and nonconceptualism of relevance to essays 2, 3 and 4. I then discuss some of the putative interconnections between the essays. Finally, I briefly explain the relationship between the essays of this dissertation to my own recent previous work on in *Content in Thought and Perception*,¹ before adding a concluding remark.

1. On the notions of perception and perceptual experience

To see a table before one, to hear a car approach, or to feel an alabaster pyramid in one's hand are all examples of perception. In general, the proper operation of our visual, auditory, tactile and other sensory-perceptual systems involves or constitutes perception. Perception, at least normally, involves perceptual or (as one might also say)

¹ Nes (2006).

sensory experience. But the notion of perception ought not to be taken to be equivalent with that of perceptual experience. One reason for this is that there is evidence to suggest that perception in exceptional cases can take place in the absence of relevant sensory experience. For example, the phenomenon of blindsight has been taken to suggest that it is possible to see something, e.g. that a bar is oriented vertically rather than horizontally, without having any conscious experience in the part of the visual field where the bar is presented.²

Even if there were no such possible cases, however, there would be reason to keep the notions apart. It is useful to retain the notion of perception as a notion of a mental capacity that may, and plausibly does, include more than conscious perceptual experience. For one thing, perception at least normally includes the acquisition of knowledge. Indeed, perception is sometimes defined as the acquisition of knowledge about certain kinds of objects by certain kinds of means. Stephen Palmer's definition in a much-used textbook is a case in point:

In the context of this book, *visual perception* will be defined as the process of acquiring knowledge about environmental objects and events by extracting information from the light they emit or reflect. (Palmer 1999: 5)

Such perceptual knowledge should not, however, be equated with perceptual experience: as we shall discuss later, knowledge is plausibly belief-dependent while perceptual experience plausibly is not. For another thing, perception involves various sorts of subconscious perceptual states or processes, which contemporary cognitive

² See Weiskrantz (1990).

science has had significant progress in investigating.³ Several of these states ought, arguably, to be counted as mental states, in that they have representational significance and various computational or broadly inferential processes are defined over them. But these subconscious perceptual states are not perceptual experiences, precisely because they are not conscious.

A vital characteristic of the notion of perceptual experiences is, then, that it is a notion of conscious states.⁴ Clearly, the notion of consciousness applied here is one of *state-consciousness*, not of *creature-consciousness*. Perceptual experiences are not only states of conscious creatures (creatures that are awake, not comatose or dreamlessly sleeping), but themselves states of consciousness. Moreover, the notion is one of phenomenal consciousness, not of access consciousness in the sense of Block (1995). There is something it is like for one to have a perceptual experience; perceptual experiences are states with a certain phenomenal character. For example, there is something it is like to look at the sunset, or to listen to rain hammering on the roof. At the same time, although the notion of consciousness in play is one of phenomenal, not of access, consciousness, it remains that perceptual experiences are things we (i.e. normal human adults) at least ordinarily have introspective access to: we can reflect on experiences in an apparently direct and intimate way and thus tell at least a few things about what they are like.⁵

³ For an overview of visual cognitive science, see Palmer (1999).

⁴ Perhaps it will be protested that the notion of perceptual experiences is also, or exclusively, a notion of conscious *events*. I do not deny that there are perfectly legitimate notions of perceptual experiences that apply either exclusively to events, or indifferently to states and events. However, for consistency, I have adopted a policy of construing experiences as states throughout.

⁵ See Botterill & Carruthers (1999) for an overview of different notions of consciousness, including the ones distinguished in this paragraph.

The state-conscious nature of perceptual experiences does not distinguish them from beliefs; at least not from beliefs of the occurrently held kind. It is a controversial matter whether their status as phenomenally conscious distinguish them from occurrent beliefs or judgements.⁶ Still, it is relatively uncontroversial that perceptual experiences are not beliefs. The persistence of illusions after they have been subjectively recognised as such vividly illustrates this. When I realise that the two lines in the Mueller-Lyer illusion are equally long, I no longer believe them to be unequally long, though my experience of them need not and typically does not change: they still look unequally long to me.⁷ Thus perceptual experiences are belief independent. Since knowledge plausibly entails belief, this also shows that perceptual experiences are not cases of (propositional) knowledge.

As anticipated, I equate perceptual experiences with sense experiences, i.e. with experiences that somehow involve or depend on the use of one or more of our sensory modalities.⁸ This means that if we have some non-sensory, perhaps somehow intellectual, perceptual or quasi-perceptual faculties, and if some essentially phenomenally conscious states are involved in the use of these faculties, those states are not perceptual experiences. This exclusion seems justified since such 'intellectual' experiences, if there are any (which in itself is a highly controversial question), would *prima facie* seem to be very unlike visual, auditory, or other sensory experiences.

⁶ See Pitt (2004) and Robinson (2005) for discussion of whether thoughts or beliefs have phenomenal character.

⁷ Armstrong (1968) and Pitcher (1970) argue, in effect, that if perceptual experiences are anything, they are tendencies to acquire beliefs. Their theories are powerfully criticised in Jackson (1977).

⁸ Note that this characterisation allows that some perceptual experiences may be essentially cross-modal in character. See Tye (2003) for an argument that some are.

I also contrast perceptual or sensory experiences with states of sensory imagination. Thus I will not count states of visual imagery, for example, as cases of visual experiences (although they seem to have a broadly visual phenomenal character). This exclusion may be motivated among other things by a salient functional feature of perceptual experiences, viz. that they are not subject to one's will. Of course, various conditions necessary for having perceptual experience of a certain kind are subject to one's will: whether one's eyes are open, whether one's hand is brought into contact with something, etc. Various factors that influence the detailed nature of one's perceptual experience are similarly under voluntary control: where one directs one's gaze, what one puts before one's sensory receptors for them to pick up on, etc. Still, it remains that perceptual experiences in an important sense are things that happen to one, not things one does. Sensory imaginings seem by contrast to be typically or standardly subject to one's will.

I take perceptual experiences to include not only the immediate, phenomenally conscious states involved in normal, 'veridical' perception, but also the apparently similar states that happen when we suffer from a sensory illusion or hallucination. This inclusion, like the exclusion above, again may be motivated by appeal to (among other things) the noted salient functional feature of perceptual experiences: experiences in sensory illusion or hallucination seem no less immediate, and disconnected from one's will, than sensory experiences in normal, 'proper' sense perception. Note that in making this inclusion, I in no way intend to prejudge whether the perceptual experiences involved in normal, proper perception and those involved in illusion or in hallucination are of the same fundamental mental kind, a question characteristically denied by

disjunctivists.⁹ The inclusion may, however, have implications for the relationship between the notions of perception and of perceptual experience. If we see perception as Palmer and others do, as something that essentially involves the acquisition of knowledge about environmental objects, by means of extracting information from physical impressions on our senses (in a relatively informationally encapsulated way), it seems that we cannot take perceptual experiences exclusively to be involved in perception.¹⁰ When we suffer from illusion or (especially) from hallucination, we are capable of acquiring knowledge of objects in our vicinity by extracting information from the relevant sensory inputs (in a relatively informationally encapsulated way): we are misled rather than informed about our environment. So on a conception of perception along Palmer's lines, cases of illusion and of hallucination are not ones of perception.¹¹

In contrast, finally, perceptual experiences with bodily sensations, such as pains, itches, and the like. This exclusion is motivated mainly on the grounds that bodily sensations traditionally have been kept apart for separate discussion in the philosophical literature, and not been treated in the same breath as visual or auditory experience, say. Again, the exclusion is not intended to prejudge the interesting question about the extent of the similarities between, for example, visual experiences and pain experiences.

⁹ For a discussion and defence of disjunctivism, see Martin (2002, 2004).

¹⁰ See Fodor (1983) for the notion of informational encapsulation invoked here. The parenthetical restriction invoking this notion seems to be needed in order for the notion of perception to be properly constrained. Without it, even highly theoretical inferences, relying on specialist knowledge and training, about the environment, based in some way on optical inputs, would count as cases of visual perception.

¹¹ While there is much to be said for a definition of perception along Palmer's lines, there is also an entrenched use of the term 'perception' on which sensory illusion and hallucination count as phenomena of sense perception. Some of my uses of the term will be of this latter, broader sense. I trust that this minor ambiguity in the term 'perception' will not cause major problems for the interpretation of the arguments to follow.

These remarks on the notion of perceptual experience do not (and are not intended to) give clear answers to all questions about the extension of the term ‘perceptual experience’. Notably, they do not give an obvious answer to whether we have perceptual experiences when we are dreaming. If the kinds of conscious states implicated in dreaming that seem to have a broadly sensory phenomenology are cases of sensory imagination, it follows from what I have said that they are not perceptual experiences. If these states are cases of hallucinations, it follows from what I have said that they are. Some awkwardness attends to either of these classifications of dream experiences: they are, at any rate, not paradigms of either sensory imaginings or of hallucinations. I must, unfortunately, leave the interesting question of their proper classification for another occasion.

The notion of perceptual experience characterised by the remarks above largely concurs, I believe, with the standard use of the term ‘perceptual experience’ in contemporary philosophy of mind. I follow the tradition in concentrating on visual experience as the paradigm instance of perceptual experience. In doing so, I also take on board the (often implicit) optimistic idea according to which much of what one finds by concentrating on this instance legitimately may be generalised to the wider class of perceptual experiences. If this optimism is unjustified, what follows can be read more restrictively as claims about visual experience.

2. Essay 1: weak representationalism

The first essay in the dissertation addresses a question that is foundational to the three essays that follow it. The question is whether perceptual experiences have

representational content. The claim that they have such content has been labelled ‘weak representationalism’ (Chalmers 2004), and has become something of an orthodoxy in philosophy of mind over the last couple of decades. In this essay, I defend weak representationalism against some recent, fundamental objections put forward by Charles Travis (2004).

To say that perceptual experiences have representational content is to say that they represent something as being the case, and that they thus allow for a non-derivative classification as veridical or non-veridical depending on whether what they represent as being the case is indeed the case. If perceptual experiences have representational content they are analogous in at least *that* respect to such paradigm propositional attitudes as beliefs or desires. Any belief, and any desire, has a representational content that determines, respectively, the conditions under which the belief is true, or the conditions under which the desire is fulfilled. Further discussion on what should be taken to be included in the very thesis of weak representationalism will be found in the early sections of essay 1.

Weak representationalism (for short, representationalism) has become popular for a number of reasons. One of the more influential is that it has seemed to provide an attractive account of illusory or hallucinatory experience. In an illusion, something may look bent to one even though it is in fact straight. In a hallucination, it may be to one just as if one is seeing a mug on a table in an ordinary room, even though one is in fact a lone brain in a vat on a far-away planet. Illusions and hallucinations involve sensory experiences, of which we must provide some sort of account. The illusory experience cannot consist in one’s seeing a straight physical thing – in one’s having a straight

physical thing visually presented to one – since no straight physical thing is seen, or even needs to be around to be seen in the case at hand. Similarly, the hallucinatory experience cannot consist in one's seeing a physical mug. It cannot indeed consist in one's seeing any relevant physical object, for no physical thing need be interacting with one's brain in the way required for one to be seeing it.¹² If visual experiences, even in illusion and hallucination, consist in seeing a particular, it seems these experiences must consist in seeing a non-physical particular. This conclusion has been affirmed by a number of philosophers, including classical empiricists such as Locke and Berkeley, and such sense-datum theorists as Russell (1912), Moore (1918-9) and Broad (1923) in the early twentieth century.

However, for a variety of metaphysical, epistemological and phenomenological reasons, philosophers have increasingly rejected the idea that illusion or hallucination consist in seeing, or being otherwise visually aware of, non-physical particulars. To many of these philosophers, the idea that perceptual experience has representational content has seemed to provide an attractive alternative. This idea allows us to say that visual experiences, even in illusion and hallucination, *in a sense* involve being aware of objects. It allows us to say this in the same sense as that in which one is aware of centaurs when thinking about centaurs, or aware of polka-dotted Rolls Royces in wanting to be the first to own one. In thinking and desiring we can, pretty uncontroversially, represent things there are not. We do not therefore need to posit non-physical centaurs, or ditto polka-dotted Rolls Royces, to account for the possibility of the thought of centaurs or the desire for a polka-dotted Rolls Royce. If perceptual

¹² I am assuming that seeing a physical object requires a suitable sort of causal interaction with it. Note that this assumption in no way entails the controversial claim that seeing can be analysed in terms of causation and mental states that fall short of seeing. For discussion of seeing and causation, see e.g. Grice (1961), Snowdon (1980-1), Lewis (1980), and Child (1994).

experience fundamentally is representational in nature, we similarly do not need to posit non-physical straight things, or non-physical particulars of other sorts, to account for the illusion and hallucination.¹³

Other reasons given for representationalism include include epistemological (McDowell 1994; Brewer 1999), phenomenological (Harman 1990; Tye 2000), semantic (Tye 2000), and cognitive-scientific arguments (Dennett 1991). I make no attempt to survey these arguments here or in essay 1. My point of departure is rather with Travis's (2004) rich, complex and fundamental criticism of the view. He argues, as I interpret him, that there are no non-question-begging facts about perceptual experiences that ground the non-arbitrary ascription of a given representational content (as opposed to countless others) to a given perceptual experience, and thus that perceptual experiences are not as such either veridical or delusive. He also argues that the fact that experiences can be misleading, as in illusions, gives no support for representationalism.

The major part of my reply is devoted to examining Travis's case that facts about how things look to one, in having a visual experience, do not determine a coherent representational content, for that experience. I argue that Travis's case here fails to respect an important distinction between comparative and non-comparative looks reports, and that he consequently leaves unharmed the claim that the facts expressed by non-comparative looks reports determine a coherent representational content. I also examine in detail Travis's suggestion that illusory experience involves seeing certain 'demonstrable looks' that things have. What, I ask, may these demonstrable looks that

¹³ This quick sketch of an argument from illusion or hallucination to representationalism skates over a number of important and controversial issues. Notably, I leave tacit the so-called 'spreading step' whereby conclusions about illusory or hallucinatory experience are generalised to veridical perceptual experiences. This step has been at the heart of recent discussions over disjunctivism.

we see be? On one answer to this question, demonstrable looks turn out to be re-branded sense data: non-physical particulars figuring as objects of visual awareness. On another, demonstrable looks turn out to be effective notational variants of representational contents. On no answer is the appeal to the seeing of demonstrable looks both independently plausible, capable of playing the role Travis wants it to play, and an alternative to a view on which perceptual experiences have representational content.

My reply to Travis's objections, however successful it might be in disarming them, does not show that representationalism is correct. Nevertheless, I believe the reply covers ground that is useful to the more positive project of vindicating representationalism. In particular, I believe it pinpoints some issues on which the representationalist would be well advised to concentrate her efforts. One of these issues is precisely the semantics and metaphysics of looks or appears talk. I hope to pursue these issues further in future work.

3. Essays 2, 3 and 4: conceptualism and nonconceptualism

If representationalism is correct, perceptual experiences are analogous to beliefs and desires at least in that they are states with representational content. The question is whether they are also analogous to beliefs and desires in any one of several other candidate respects. One respect that has been the subject of much discussion over the last couple of decades is that of conceptualisation: are perceptual experiences somehow conceptual, or conceptualised, as beliefs and desires presumably are? This discussion has been conducted notably in terms of a debate on whether perceptual experience has conceptual content, or nonconceptual content, or both. Philosophers for whom perceptual experience has conceptual content, and does not also have nonconceptual

content, are standardly identified as ‘conceptualists’, while those for whom it also or only has nonconceptual content standardly are labelled ‘nonconceptualists’.^{14, 15} The label ‘nonconceptualist’ may also be extended to embrace those who deny that perceptual experience has representational content at all (this has the attractive consequence of making ‘nonconceptualism’ equivalent to the denial of ‘conceptualism’). However, the majority of recent nonconceptualists have accepted representationalism, and posited some sort of nonconceptual representational content. Essays 2, 3 and 4 engage with this debate between conceptualists and nonconceptualists.

As noted, the debate has taken place against the background of some shared presumptions about beliefs and desires. Most importantly, it has been presumed that these states are in some way conceptual, or conceptualised.¹⁶ The notion of conceptual content has tacitly or explicitly been introduced and motivated against the backdrop of

¹⁴ The term ‘non-conceptual content’ is introduced in Evans (1982), who argues that perception and sub-personal information-processing states have content so described. Peacocke (1983) early on took a conceptualist view of experience, but later (1986, 1989, 1992, 1994, 1998, 2001) became one of the most vigorous proponents of nonconceptualism. Another influential nonconceptualist is Tye (1995, 2000, 2006a, 2006b). Notable defenders of conceptualism include McDowell (1994, 1998a, 1998b) and Brewer (1999, 2005). (I should note that Brewer no longer endorses conceptualism, and now (2006, 2007, forthcoming) defends a view on which perceptual experiences does not have representational content of either a conceptual or nonconceptual type.)

¹⁵ Henceforth, when I write that beliefs (/desires/experiences/...) have conceptual content, I mean that they exclusively have conceptual content, i.e. that they do not also have nonconceptual content, except where otherwise indicated.

¹⁶ Beliefs and desires have been invoked here principally as paradigms of higher-order, cognitive states, and of propositional attitudes. If one held that the notion of judgement, say, is better suited to such a paradigmatic role, as Evans (1982) argues, one might instead have couched the discussion in terms of a comparison of experiences and judgements. I will not rest anything on the distinction between beliefs and judgements (whatever it may be precisely), and talk primarily in terms of beliefs. I will sometimes speak in terms of ‘thoughts’, taking this to express a determinable mental state whose determinates include beliefs, judgements, suppositions, mere leisurely entertainings of propositions, etc. My use of ‘thought’ should thus be contrasted with the use of it to denote the abstract contents of states of thinking. When I write, as I occasionally do below, of ‘thought or belief’, this phrase is comparable to ‘fish or herring’ and has the same stylistic motivation as the latter.

this presumption. However, while this presumption has been shared, there is in the literature two interestingly different approaches to the notion of conceptual content. The debate has come to self-consciousness about this difference only fairly recently, over the last six or seven years.¹⁷

On the one approach, conceptual content is identified as representational content of conceptualised states of mind, of which beliefs have been held to be paradigmatic examples. Below, we shall return to the question what it is for a state of mind to be conceptualised. The notion of conceptual content has, moreover, been introduced in a way that leaves open, at least as a coherent possibility, that the very same representational content that is the content of a conceptualised state of mind, and for that reason counts as conceptual content relative to it, could also be the content of a representational state that is not conceptualised, and for that reason count as nonconceptual content relative to the latter. On this approach, the primary question is whether states of mind of a given kind are conceptualised or not; the notion of (non)conceptual content may be applied in a relative way to states once the primary question has been settled.¹⁸ The claim that perceptual experiences have conceptual content, as understood on this first approach, thus amounts to the claim that they are conceptualised states of mind. This claim has been labelled ‘state conceptualism’ (Byrne 2005) or ‘relative conceptualism’ (Speaks 2005), and its denial ‘state nonconceptualism’ or ‘relative nonconceptualism’. I will adopt Byrne’s labels. Essays 2 and 3 below critically examine state conceptualism.

¹⁷ My account of these two approaches relies heavily on Byrne (2005) and Speaks (2005). Byrne and Speaks both credit earlier work by Stalnaker (1998) and Heck (2000).

¹⁸ Crane (1992) is often identified as evincing this first approach to the notion of conceptual content.

The second approach initially identifies conceptual content as representational content that meets a certain condition, and then says that a state of mind has conceptual content just in case the content it has is identical to a representational content that meets the condition. Unlike the first approach, this approach does not leave it open as so much as a coherent possibility that one and the same content could show up as a conceptual content in one state of mind and as a nonconceptual content in another. We may distinguish two ways of implementing this second approach, depending on how conceptual content initially is specified. The first implementation specifies conceptual contents as representational contents with a certain intrinsic property. Typically, the claim is that conceptual contents are built out of concepts, or out of ‘modes of presentation’, that satisfy a certain principle of individuation, such as ‘Frege’s Principle’ that content ...*F*... and content ...*G*... (obtainable from ...*F*... by substituting *G* for *F* at one or more places) are distinct if it is possible for someone rationally to believe ...*F*... while doubting or denying ...*G*....¹⁹ On this first implementation, it becomes a substantive and indeed controversial claim that beliefs or thoughts have conceptual content. For example, if a neo-Russellian approach to content of belief of the sort advocated by a Salmon (1986), Soames (1987) or Thau (2002) turns out to be right, or an ‘ultra-coarse-grained’ approach of the kind defended by Stalnaker (1984) or Lewis (1983) does, beliefs and thought do not have conceptual content. On the alternative, second implementation, the claim that beliefs and thoughts have conceptual content is, by contrast, taken as stipulative, in that the category of conceptual contents initially is identified as representational contents of the kind beliefs and thoughts have.²⁰ More precisely, conceptual content may be identified as representational content that is the

¹⁹ Peacocke (1992, 2001) defines conceptual content as contents composed out of modes of presentation the individuation of which accords with Frege’s Principle.

²⁰ See Byrne (2005) for a characterisation of conceptual content of this kind.

content of a possible thought or belief. Thus, the claim that experiences have conceptual content, on this second implementation, is to be understood as the claim that each perceptual experience has a content that is the content of a possible thought or belief. In this dissertation, I will restrict attention to the second way of implementing this second approach. Following Byrne (2005), I will label the claim that perceptual experiences have conceptual content, understood in this way, ‘content conceptualism’, and the denial of that claim ‘content nonconceptualism’. Essay 4 below critically examines an issue concerning content conceptualism.

As noted, it is only quite recently that the debate has come to self-consciousness that there are these two approaches to the notion of conceptual content in it. This is probably because it has tacitly been assumed that they are equivalent, i.e., that state conceptualism is correct iff content conceptualism (on either of the two implementations) is. In sections 7 and 8 below, I shall return to the relation between state and content conceptualism.

In the remainder of this section, I first seek to make a little clearer the notion of a conceptualised state presumed by state conceptualism. I then give an overview of the arguments of essays 2, 3 and 4.

3.1 Conceptualised states

To get at the notion of a conceptualised state, it is useful to start from an apparent feature of beliefs and desires. It seems that someone can believe that cats are mammals, for example, only if they have a concept of mammals. Similarly, it seems that someone can desire a glass of soda, only if they have a concept of soda. Arguably, a constraint on beliefs and desires along these lines is enshrined in folk psychological practice. We are

willing to attribute to a three-year-old the belief that a cup, if unsupported, would fall to the ground, but we are typically not willing to attribute to her the belief that gravity acts on the cup. If asked about this difference, a compelling answer, from the point of view of common-sense psychology, is that our three-year-old has no concept of gravity. Many philosophical arguments also tacitly assume a constraint on propositional attitudes along these lines. For example, a familiar objection to the KK principle in epistemology – the principle that someone knows only if they are in a position to know that they know – is that someone can know something, say that water is wet, without having the concept of knowledge.

It is a non-trivial matter to spell out in a general way the constraint on beliefs and desires of which we have just seen some putative instances. A first stab at it is the following:

Conceptualisation constraint – belief (v.1) A subject can believe that ...F..., only if she has a concept of F.

This formulation seems to point in the right direction – to convey the general idea. However, as it stands, it is afflicted by the problem that it is ill-formed for some instances of ‘F’ (I am assuming that ‘F’ holds the place for a semantically unitary expression). The formulation ‘A subject can believe that snow is white and grass is green only if she has a concept of and’ is simply ungrammatical.

One way in which one might seek to remedy this problem rests on the assumption that each semantically unitary expression ‘F’ used in specifying the content of a belief has a

semantic value, in the context in which it is thus used. Using this idea, the constraint is rephrased in meta-linguistic terms:

Conceptualisation constraint – belief (v.2) A subject can be such that ‘x believes that ...F...’, as uttered in context C, is true of her, only if she has a concept of the thing which is the semantic value of ‘F’, as uttered in C.

‘Thing’ is to be interpreted liberally here as embracing any object, property, function or whatever. To take an example, suppose I have a conversation with my friend Will about a car he has recently come into possession of, and say:

(1) Mike thinks you stole that [nodding at the car] yesterday and grass is green.

The semantic values of ‘you’, ‘that’, and ‘yesterday’, in the context in which (1) was uttered, are, respectively, Will, the car Will has recently come into possession of, and (let’s say) 10.07.2007. The semantic values of ‘stole’, ‘and’, ‘grass’ and ‘green’ are, I will presume, respectively, the relation of stealing, the truth-function conjunction, the kind grass, and the colour green. Thus (v.2) tells us that (1) can say something true of Mike only if he has concepts of Will, the car Will has recently come into possession of, the day 10.07.2005, the relation of stealing, the truth-function conjunction, the kind grass, and the colour green. Note that this is not to be read as saying that Mike has a concept of the semantic value of ‘that’, in the relevant context, only if he has the descriptive concept *the semantic value of ‘that’ in C* or the descriptive concept *the car Will has recently come into possession of*. On the intended reading of (v.2), it is consistent with (v.2) that (1) is true even though Mike no notion of semantic values as such nor has ever heard of Will but only seen him when he appeared to be hot-wiring a

car. (v.2) does not lay any restrictions on the kinds of concepts of objects, or of properties, one needs to have to have a certain belief about that object or property.

We may however seek to get the same benefits as (v.1) in simpler, non-meta-linguistic terms, if we allow ourselves some flexibility in the schematic formulation.

Conceptualisation constraint – belief (v.3) A subject can believe that ...F... only if she has a concept of F/F-ness/F-hood/...

The slashes at the end of this schema are to allow that different instances of it will take a slightly different form, depending on the semantic type of the expression taking the place of 'F'. The dots are there to allow ourselves some open-endedness as to how an appropriate term for the semantic value of 'F' is to be formulated.

Each of the versions of a conceptualisation constraint formulated so far employs the notion of having a concept of an object, or a property, or a truth-function, or whatever. How is this notion to be understood? As far as possible, I seek to remain neutral in this dissertation between various detailed philosophical theories of concepts.²¹ I will operate on the basis of the following relatively uncontroversial platitudes about concepts. First, concepts are things with semantic properties: a concept is a concept of something or other. Second, concepts are things with combinatorial properties that are analogous to those of words. We use a given concept in combination with other concepts to form mental states with a propositional content that reflects the combination of concepts used, in a manner that is analogous to the way in which we use a given expression in

²¹ For an overview of different philosophical theories of concepts, see Margolis & Laurence (1999).

combination with other expressions to form utterances with a propositional content that reflects the combination of words used. A third and closely related point is that concepts can be classed into certain semantic types that register their semantico-combinatorial properties, just as words can: just as it makes sense to talk of singular terms, predicative terms, and quantificational terms, it makes sense to talk of singular concepts, predicative concepts, and quantificational concepts. Forth, concepts are many-one to the objects or properties or whatever they are concepts of. There can be many different concepts of a given object or property. An amnesiac who wakes up in Inverness and can think of it only as this town he now finds himself in, has a different concept of the town from me, who has only heard of a town called 'Inverness'. Yet my and the amnesiac's concepts are concepts of the same town. Fifth, and obviously, concepts are things subjects can have or possess. For a thinker to possess a concept is for her to have a representational capacity of a certain sort: a conceptual representational capacity. We shall return to what it is to possess a concept in essay 3 below where I explore in some detail how conceptual and nonconceptual representational capacities may be distinguished from each other.

The fourth platitude about concepts – that they are many-one to the objects or properties or whatever they are concepts of – might suggest to some readers a weakness of the three versions of conceptualisation constraint formulated above. None of these lays any constraints on what kind of concept of a given object or property etc. one needs to possess to have a given belief. Yet one might think that a proper conceptualisation constraint on belief should incorporate such a restriction. One way of attempting to do so is by means of the assumption that each semantically unitary expression 'F', as used in a context C of specifying the content of a given belief, expresses a concept F_c in that

context *C*. As we did in (v.2), this assumption is easiest to incorporate on a meta-linguistic formulation of the constraint:

Conceptualisation constraint – belief (v.4) A subject can be such that ‘*x* believes that ...*F*...’, as uttered in context *C*, is true of her, only if she has the concept F_c , expressed by ‘*F*’ in *C*.

However, also this formulation confronts some problems. For one thing, one might doubt whether each semantically unitary expression expresses a concept, even in a given context where it is used to specify the content of a certain belief. In particular, it is controversial whether proper names do so. Even setting this aside, the formulation does not seem to capture in all cases the intuitive underlying constraint on beliefs that we want to get at. Suppose Jimmy and I attend a meeting, and that a perspicuously sweating and panting man enters at one point. Jimmy discreetly nods towards him and says:

(2) That guy has been in a hurry.

As it happens, the fellow Jimmy refers to is the one my colleague Mike and I know as Sam Liggins. Later I remark to Mike:

(3) Jimmy believes Sam Liggins has been in a hurry.

On the face of it, I am speaking truly in uttering (3).²² If we apply (v.4) to (3), we get the consequence that Jimmy must possess the concept *Sam Liggins* which (we assume) ‘Sam Liggins’ expresses in the context where (3) is uttered. However, if Jimmy has never so much as been introduced to the name, it seems that he does not or anyway need

²² See Richard (1990) for rich and interesting discussions of the truth-conditions of belief reports. Richard allows for some contexts of use on which (3) may not be true, but holds that (3) typically is true in ordinary contexts of conversation.

not possess this sort of proper-name-associated concept of the man who is in fact Sam Liggins. In order for me to utter (3) truly, it would seem to suffice for Mike to have some sort of visual, demonstrative concept of Sam Liggins.

A notion invoked by some writers on conceptual and nonconceptual content of use in dealing (perhaps among other things) with this sort of worry is that of a ‘canonical specification’ of the content of a mental state.²³ The idea is, roughly, that the conceptualisation constraint tells us what concepts one needs to possess to have a given belief only on the basis of a canonical specification of its content. The intended notion of a canonical specification should exclude (3), for example, as not giving a canonical specification of the content of the belief ascribed to Jimmy. On the other hand, it should presumably include the self-ascription Jimmy would give of his belief:

(4) I believe that that guy has been in a hurry.

It is not straightforward to spell out in general terms what it is for something to be a canonical specification of a content. Adrian Cussins says that content is ‘canonically characterised by a specification which reveals the way in which it presents the world’ (1990: 383 fn. 25). However, the phrase ‘the way in which it presents the world’ is ambiguous. One interpretation of it starts from the (fairly standard) characterisation of the content of a belief (or desire, or hope, or other representational state) as the way the world is believed (or desired, or hoped, or otherwise represented) to be. On this interpretation, a canonical characterisation of the content of the belief is a characterisation that specifies how the world is believed to be, in that belief. Now, the way the world is believed to be, in believing that that guy over there (intending Sam Liggins) was in a hurry, is just the way it is believed to be, in believing that Sam

²³ See Cussins (1990) and Crane (1992). Crane and Cussins also intend the notion of a canonical characterization to do other work.

Liggins was in a hurry.²⁴ Thus, on this first interpretation of ‘the way in which it presents the world’, we do not obtain the desired contrast between (3) and (4) in point of canonical specification of content. A second, and more promising, interpretation of the phrase ‘the way in which it presents the world’ takes its departure from the Fregean or neo-Fregean notion of a ‘mode of presentation’.²⁵ In terms of this notion, we may distinguish two beliefs in the vicinity of (3) and (4). Both of these are beliefs about Sam Liggins, to the effect that he was in a hurry. They differ however with respect to the mode of presentation under which Sam Liggins is thought of. In one, Sam is thought of via a visual, demonstrative mode of presentation. In the other, he is thought of by a mode of presentation that (we assume) is somehow indicated or invoked by the name ‘Sam Liggins’ (perhaps as used in a given context). In terms of this distinction, one might posit that (4) provides a canonical characterisation of the content of the first but not the second of these beliefs, and vice versa for (3). A canonical characterisation of the content of a belief is, then, one that is sensitive not only to its truth-conditions, and indeed not only to the objects referred to or the properties attributed to them, but also to the modes of presentations of the objects and properties. Adopting the notion of canonical specification, so understood, we may formulate the following constraint:

Conceptualisation constraint – belief (v.5) A subject can be such that ‘x believes that ...F...’, as uttered in context C, where this gives a canonical specification of content of the belief attributed, is true of her, only if she has the concept F_c , expressed by ‘F’ in C.

²⁴ In making this claim, I am assuming that a description theory of proper names and demonstratives is false. Note that this is not tantamount to assuming that a Fregean theory of proper names and demonstrative is false, since there are broadly Fregean (or neo-Fregean) theories of proper names and demonstratives on which they are not descriptive, e.g. Evans’s (1982).

²⁵ See Frege (1892). Neo-Fregean views are developed in Evans (1982), Peacocke (1992), and McDowell (1998c).

For each of the versions of the conceptualisation constraint considered so far, the supposition is that they hold *pari passu* for desires and other central propositional attitudes such as hoping that P. For each, one can query whether it generalises to perceptual experience. The claim that contentful states of mind of a certain sort are conceptualised just is the claim that a generalisation of the conceptualisation constraint to states of that sort holds good.

In what follows, I will be more interested in whether or not a subject who has an experience needs to have a concept of the properties she experiences things as having than whether she needs to have a certain kind of concept of these properties as opposed to other kinds of concepts of the same. Thus, for our purposes, the greater refinement of versions 4 and 5 in general will not matter. Since version 1 was seen to be problematic, and since 3 has the advantages over 2 of being simpler and non-meta-linguistic, version 3 will be my canonical conceptualisation constraint on belief in what follows. The question I will be examining in essays 2 and 3 is whether version 3 of the conceptualisation constraint generalises to perceptual experience, i.e. whether a subject can have a perceptual experience as of ...F..., only if she has a concept of F/F-ness/F-hood/...

3.2 Essay 2: conceptualism, recognitional concepts and indiscriminability

Essay 2 begins with a putative counterexample to the conceptualisation constraint on experience. Visual experiences are extremely specific, or fine-grained, in the colours, or shapes, or locations, we can experience things as having. For example, on some estimates we can sensibly discriminate more than ten million separate shades of

colour.²⁶ Yet one might think that we couldn't really possess a concept for each of these myriad shades (or shapes, or other fine-grained properties) we experience things as having.²⁷

An influential conceptualist response to this problem is McDowell's (1994). He argues that even though we do not in general have pre-existing, linguistically articulated concepts of fine-grained shades of colour prior to experiencing something as having them, the very experiencing of something as having a fine-grained shade makes it the case that one comes to be in possession of a concept of it. The concept one thus comes to be in possession of, McDowell argues, rests on a certain minimal recognitional capacity: a capacity to recognise other things as having the same shade, when one is suitably presented with them.

Several nonconceptualists have attacked this suggestion of McDowell's on the grounds that it places unrealistic demands on our capacities of sensory memory. As an empirical matter, we are not capable reliably of re-identifying as determinate shades of colour across time as we are capable of discriminating at a time. Our recognitional capacities are thus more coarse-grained than our experience. Thus, if the most determinate concepts of colours, shapes, etc. that visual experience provides us with rest on these recognitional capacities, these concepts will not be as fine-grained as the experience,

²⁶ See Tye (1995).

²⁷ Gareth Evans (1982) filed this objection early on. Presumably quite independently, Kelley (1980) in effect presses the same objection, directing it against Armstrong's and Pitcher's belief analyses of perception.

and the appeal to recognitional concepts does not solve the conceptualist's problem of fineness of grain.²⁸

In essay 2 I develop an alternative argument for the same conclusion. Unlike the above argument, the case is independent of an appeal to limitations of sensory memory. (I contend that this is an advantage of the argument, in that it refutes, if successful, even a weaker, 'cross-modal' as opposed to 'cross-temporal', version of the recognisability requirement on recognitional concepts.) Instead, the argument appeals to the non-transitive structure of phenomenal indiscriminability, or matching. I argue, first, that this non-transitivity grounds a *prima facie* sceptical challenge to the possibility of knowledgeable perceptual re-identification of fine-grained shades of colour. I then examine McDowell's and Brewer's (1999) closely related conceptualist treatments of the non-transitivity of matching. These treatments are found to ground a plausible-seeming reply to the sceptical challenge, but to give rise to an independent fatal problem over the possibility of perceiving fine-grained changes in colour. I argue that the revisions the latter problem calls for strongly re-invigorate the sceptical challenge towards knowledgeable perceptual re-identification of fine-grained shades.

3.3 Essay 3: Conceptualism, self-conscious thought and systematicity

Essay 3 continues the critical examination of state conceptualism (the claim that the conceptualisation constraint on experience holds good), building on the results of essay 2. The conclusion I aim to establish in essay 3 is that conceptualists are committed to giving up two principles about concept possession that have been widely granted, not

²⁸ Arguments of this kind are developed in Raffman (1995), Tye (1995), Kelly (2001) and Peacocke (2001).

least by conceptualists themselves.²⁹ I call these the Kantian Principle and the Systematicity Principle. The former says that in order to possess a concept one must be capable of using it in thinking self-conscious thoughts (thoughts available for immediate self-ascription); the latter that in order to possess a concept one must be capable of using it in having mental states that exhibit a certain systematicity or re-combinability. I do not argue that the commitment to deny these principles refutes conceptualism, since I do not argue that the principles are non-negotiable. I maintain, however, that abandoning them comes at least at a dialectical cost. The cost is that the conceptualist either must argue that not even thoughts exhibit the property of systematicity, or is left in the embarrassing position of being at a loss to say what the principled distinction may be between conceptual and nonconceptual representational capacities, between possession of concepts and having a representational capacity that does not constitute the possession of concepts.

The argument to this conclusion proceeds in three stages. The first stage begins with the problem of fineness of grain. I observe that the conceptualist and his nonconceptualist opponent seem to have been able to agree that the conceptualist only has three kinds of candidate concepts available to him to solve this problem: lexical concepts, recognitional concepts, and demonstrative concepts. (The classes are not necessarily mutually exclusive.) It is pretty uncontroversial that concepts of the first kind will not do the job for the conceptualist. I take essay 2 to show that concepts of the second kind will not do it either. In this section, I argue that the same goes for concepts of the third kind. The argument is, quite simply, that possession of demonstrative concepts

²⁹ For the rest of this sub-section, 'conceptualism' means 'state conceptualism'.

presumes attention to the object or property the concept is a demonstrative concept of, whereas perceptual experience of an object or property does not.

At the next stage, I ask why one should think (as both sides of the debate appear to have) that the three kinds of concept distinguished – lexical, recognitional, demonstrative – exhaust the conceptualist's candidates. While there are various themes in the literature that more or less clearly seem to bear on this question, there has been little explicit argument to the conclusion that the three classes of concepts exhaust the options. In this section, I try to remedy this, articulating two lines of argument to this conclusion. The Kantian Principle is found to be the shared central plank of these arguments, and to be the only premise the conceptualist reasonably can deny in order to avoid their conclusion (as he must if his position is not to stand refuted).

At the third stage, I ask what the principled distinction may be between conceptual and nonconceptual representational capacities, a question to which both the conceptualist and his opponent ought to have an answer. In giving up the Kantian Principle, the conceptualist ipso facto loses one resource for answering it. What alternatives are available to him? It is natural to turn here to the idea that conceptual cognition is somehow essentially systematic, or meets what Evans called the 'Generality Constraint'. This idea underlies the Systematicity Principle, according to which (to remind ourselves conceptual capacities essentially are exercisable in mental states exhibiting systematicity. I argue however, that the considerations that force the conceptualist to give up the Kantian Principle have the implication of his being prevented from affirming the Systematicity Principle. The reason is that these considerations have the implication that the conceptual capacities that (according to the

conceptualist) are operative in perceptual experience, and endow it with its content, *only* are exercisable in having perceptual experiences; yet perceptual experiences are not, as I show by means of a range of examples, states of mind that exhibit systematicity.

3.4 Essay 4: content conceptualism and diaphanousness

Essay 4 shifts focus from state to content conceptualism. To remind ourselves, this is the thesis that the content of a perceptual experience is the content of a possible belief or thought. In this essay, I argue that this thesis, on certain independently plausible assumptions, is incompatible with a much-discussed and relatively popular thesis of diaphanousness, namely this:

Thesis of Diaphanousness, (D) In introspection, one is or becomes aware of the phenomenal character of visual experience only *by* attending to apparent objects of experience and properties they appear to have.³⁰

This thesis is but one of the claims and ideas regularly mooted in the discussion of the diaphanousness of experience.³¹ I do not maintain that no other claim from this nexus of considerations has equal title to be singled out as *the* thesis of diaphanousness. My claim is only that (D) is not obviously stronger and less popular than the many other, more or less closely related, claims of diaphanousness that have been so much discussed in recent philosophy of consciousness. Nor is my aim in the paper to support the claim that (D) is true. My aim is only to argue that (D) and content conceptualism are incompatible, without assigning blame to either.

³⁰ The formulation of this thesis is inspired principally by Stoljar (2004).

³¹ The general idea of diaphanousness (or transparency, as it is also called) goes back at least to Moore, who famously asserted that 'when we try to introspect the sensation of blue, all we can see is the blue; the other element is as it were diaphanous.' (1903: 25)

The argument is this. If content conceptualism is true, then for any perceptual experience there is a thought with the same content as it. Presumably, one could have the thought without at the same time having the experience, and so without at the same time being in a state with the phenomenal character distinctive of the experience. Suppose, then, that one first has merely the thought and then immediately after a perceptual experience with the same content. The onset of the experience should involve a considerable, introspectively detectable shift in the phenomenal character of one's state of mind. How is this shift to be detected? Given (D), and the thesis of representationalism (that content conceptualists are committed to), one is or becomes aware of the new phenomenal character only by attending to the representational content of the experience. Yet, in so far as the representational content of one's states of mind are concerned, coming to have the experience *ex hypothesi* involves no change. So if (D) is true, it seems that coming to have the experience should involve no introspectible change in the phenomenal character of one's state of mind, contrary to hypothesis.

In essay 4, I set out this argument in greater detail, respond to some objections, and consider what options it leaves the content conceptualist with. I argue that she has no plausible way round denying (D).

4. Essay 5: Phenomenal concepts in introspection and perception

An interesting feature of the phenomenal character of perceptual experience is the apparently special characteristics of our knowledge of it. To know the phenomenal character of a given perceptual experience is to know what it is like for someone to have that experience. It seems that it is possible to know all the physical facts about the world

and yet not have this knowledge. Frank Jackson's (1982, 1986) example of the super-scientist Mary, brought up and trained to physical omniscience in an exclusively black-and-white environment, is designed to illustrate this possibility. According to the story, Mary knows all the physical facts about the world and yet does not know everything, since she does not know what it is like to see something red. Jackson argued that this possibility refutes the physicalist claim that all facts about the world are physical facts – that it shows that the phenomenal character of an experience is not a physical fact about it.

To know what it is like to have a certain experience is to have phenomenal knowledge of a broadly introspective kind. A number of theorists have argued, *pace* Jackson, that the special characteristics of phenomenal knowledge are not necessarily a reflection of a special non-physical nature on the part of the phenomenal characters known about, but rather of distinctive features of the *concepts* involved in having such knowledge. Such knowledge, they argue, involves the use of certain 'phenomenal concepts', grasp of which presumes one has had certain correlated experiences, or at least can imagine them, and thus is in a position to know what these experiences are like. The special characteristics of these concepts are consistent with their *referents* themselves being physical properties. At any rate, they are consistent with their referents being such that one could have 'book learning' (knowledge acquired by black-and-white means, though not necessarily physical knowledge) of the distribution of these properties. This account of what goes on in Jackson's knowledge argument I refer to as a 'perspectivalist' response.

Jackson's knowledge argument is an 'experience-directed' knowledge argument. The critical knowledge for which complete book learning *ex hypothesi* does not suffice is knowledge about *experiences*. In essay 5, I consider an analogous 'object-directed' knowledge argument, where the critical knowledge for which complete book learning *ex hypothesi* does not suffice is knowledge about the *objects* of experience: in particular, knowledge of what the objects of visual experiences are like, colour-wise. I argue that if a perspectivalist response to Jackson's experience-directed argument is correct, as I believe is plausible, then a parallel perspectivalist response to our object-directed knowledge argument is correct. This response has the consequence that close analogues of the phenomenal concepts found in introspective judgement about what it is like to have a certain experiences are found in basic perceptual judgement about what certain objects of experiences are like, colour-wise or perhaps in other respects. The latter concepts may be classified as 'perceptual-phenomenal' concepts, by analogy with the 'introspective-phenomenal' concepts found in introspection. The analogies between the perceptual and the introspective case identified in this essay are relevant, I suggest, to the currently much-contested question of how we are to conceive of phenomenal concepts.

5. Essay 6: Intentionality and the mark of the mental

Perceptual experiences are but one example of a mental phenomenon. Others include bodily sensations, beliefs, desires, plans, hopes, fears, intentions, emotions: the list could be extended. Is there any unity to this apparently disparate collection of things mental? If there is none, one might worry that it calls into question the claim of psychology to have unified domain of inquiry, and be considered a unitary scientific discipline. At any rate, Franz Brentano (1874) was exercised by this problem, as were many of his contemporaries, who were intrigued by the emerging science of

psychology. Brentano's great contribution to that debate was his revival of the notion of intentionality, and his proposal that it holds the key to the unity of the mental domain. Restricted to mental *states*, his proposal entails a thesis I will single out as Brentano's thesis:

Brentano's thesis All and only mental states are intentional states.

My intention in this essay is not to engage in historical scholarship on Brentano's rich but complex discussion of the notion of intentionality. I start, rather, from a point in time much closer to our own, with Tim Crane's (1998, 2001) recent articulation of the notion. I choose Crane not only because his articulation is lucid and broadly representative of how the notion tends to be understood in contemporary philosophy of mind, but also because he explicitly argues that the notion so articulated vindicates Brentano's thesis.

Most recent discussion of Brentano's thesis has focused on the necessity of intentionality to mentality, seeking to bolster or rebut the objection that bodily sensations or moods are non-intentional. I concentrate rather on the question of sufficiency. I argue, first, that on a simpleminded but natural reading of some central passages of Crane's articulation, the non-mental state of attracting an iron ball comes out as intentional. I then consider various ways of remedying this problem. One is to explain intentionality in overtly mental terms. This unacceptably trivialises one half of Brentano's thesis. The other is to ratchet up the necessary condition for intentionality, by appealing to such features of our talk about the intentional as failure of substitution and of existential presupposition. I argue that these strengthenings of the condition

necessary for intentionality either fail to exclude all offending non-mental states, or are so strong as to make it highly doubtful whether all mental states meet it. In particular, I argue that it is doubtful whether basic sensory or perceptual states meet the conditions strong enough to exclude all non-mental states.

6. Interconnections³²

The six essays collected in the dissertation are obviously thematically connected in that they all discuss aspects of perceptual experience or perceptual judgement. There is also a number of more specific links between them, some of which are pointed out in footnotes in the essays themselves. In this section, I outline how some of the principal conclusions of the individual essays bear the issues, conclusions and arguments of the others. Of course, no attempt is made to chart these interconnections in complete and comprehensive detail. The following addresses what I take to be the most important links – or putative links – between the essays.

6.1 The bearing of weak representationalism

Essay 1 defends the claim that perceptual experiences have representational content, what I call weak representationalism, or for short, representationalism. I operate on the assumption of weak representationalism in essay 2, 3 and 4. The thesis also has an important bearing on essay 6: if perceptual experiences have representational content, they are intentional in a fairly robust sense of that term, and thus conform to the claim that all mental states are intentional even in a fairly robust sense of that term (however, and this is something to which I shall return, it does not follow from representationalism that they are intentional in the very strong sense that seems to be needed to exclude all non-mental states from the intentional).

³² It is strongly advisable to read this section and the next after having read the essays.

In operating on the assumption of representationalism in the essays that connect with the conceptualism v. nonconceptualism debates (essays 2, 3 and 4), I follow the tradition. These debates have generally been conducted against the shared backdrop of representationalism. The presumption has been that perceptual experience has representational content; the question the debate has revolved around is whether this content is conceptual or nonconceptual. One might wonder, however, whether the tradition has been correct in doing so. Could we sensibly distinguish between a conceptualist and a nonconceptualist view of perceptual experience even if representationalism is false?

Now, above we drew a distinction between *state* and *content* (non)conceptualism. We defined content conceptualism as the claim that each perceptual experience has content, where its content is identical to the content of possible belief or thought. The content of a belief or thought is representational content. Thus content conceptualism immediately entails and presupposes representationalism.

We defined state conceptualism, by contrast, as the claim that perceptual experience is a conceptualised state of mind, i.e. that it meets a conceptualisation constraint, of which we took our preferred formulation to be this,

Conceptualisation constraint – experience A subject can have a perceptual experience as of ...F..., only if she has a concept of F/F-ness/F-hood/...

Now, as I set out above, the natural way of explaining and motivating the question of state conceptualism is to note the plausibility of a certain conceptualisation constraint on belief; observe that this constraint plausibly generalises to a range of other intentional states, such as desires, intentions, hopes, mere suppositions, and so on; and then suggest that the constraint ought equally to generalise to perceptual experience. In order to motivate and indeed make sense of this generalisation, one needs, it seems, to presuppose that perceptual experiences are intentional states: that they have content in some sense of that term. After all, the requirement of conceptualisation concerns precisely the objects or properties that somehow figure in, or are determined by, the content of the state. However, it is not clear whether one strictly speaking needs to presuppose that perceptual experiences have *representational* content, in the sense of the term adopted in essay 1 and explained in outline above. It is critical to that sense that representational content determines veridicality-conditions for the experience. Veridicality-conditions are a species of truth-conditions. Thus, it critical to that sense that representational content determines truth-conditions. Now, content that determines truth-conditions, and accordingly is truth-evaluable content, qualifies as *propositional* content on a thin but recognisable sense of that term.³³ However, as I observe at the beginning of essay 3, it is controversial whether all intentional states have propositional content. Such states as admiring ballet dancers, loving a great cook, and imagining a bearded dog certainly seem to be intentional states. It is controversial, however, whether

³³ A somewhat richer sense of the term takes it to be essential to propositional content that it is content of the sort that beliefs or thoughts have: propositional contents are essentially things that are believed or may be believed. This richer sense is probably assumed when Peacocke (1992) labels a certain type of perceptual content ‘proto-propositional’ content, even though the type of contents in question clearly is truth-evaluable and determines truth-conditions. I here operate with the weaker notion of propositional content.

they are states with propositional content.³⁴ It seems we may deny that they are without denying that they are intentional states. Moreover, even if they do not have propositional content we can sensibly generalise a conceptualisation constraint to them, as in,

Conceptualisation constraint – admiration In order to admire F, a subject needs to possess a concept of F.

Indeed, this constraint not only makes sense but seems plausible.³⁵ It follows that the generalisation of the conceptualisation constraint to perceptual experience does not presuppose them to have propositional content. In other words, state conceptualism does not strictly presuppose representationalism, but only the weaker claim that perceptual experiences are intentional.

This is not to deny the relevance of representationalism to the motivation for state conceptualism. If perceptual experiences have representational content, and so, in our weak sense, propositional content, they are to that extent more analogous to beliefs than if they hadn't had such content. Since beliefs are the paradigm conceptualised states, the stronger the analogy of perceptual experiences with beliefs the stronger the prima facie reason to think they are conceptualised.

6.2 The bearing of state nonconceptualism

Essays 2 and 3 do not purport to establish the denial of state conceptualism, state nonconceptualism. Yet it is fair to say that the drift of the essays is towards this

³⁴ See Searle (1983: 6--7), Crane (2001: 112--4), and Forbes (2006: 52--68) for arguments that not all intentional states are states with propositional content.

³⁵ See footnote 2 in essay 3 for a reply to an objection to the constraint. [[XX]]

position. What would the truth of state nonconceptualism imply for the claims and arguments of the other essays in the dissertation?

Some theorists might worry that it is inconsistent with the representationalism defended in essay 1. Representationalism says that perceptual experiences are like beliefs in having representational content. What can account for this contentfulness on the part of experiences? In the case of beliefs, one might think their contentfulness rests on the fact that beliefs involve the use of concepts. If experiences do not, how can it be that they are contentful?

While we certainly need to be able to answer this question, I do not think we should be particularly worried that we might prove to lack the resources for doing so. The crucial resource we need here is the idea of a nonconceptual representational capacity. We need to be in a position non-arbitrarily to delimit a non-empty class of nonconceptual representational capacities. Essay 3 gives grounds for optimism that the state nonconceptualist can do this. We there argue that the state *conceptualist* is committed to give up two principles for concepts that have been widely assumed: the Kantian and Systematicity Principles. However, the state *nonconceptualist* is not so committed. He is free to use them to mark off conceptual from nonconceptual representational capacities. If the argument of essay 3 is sound, the state conceptualist cannot come back with the objection that the notion of a representational capacity that does not meet these principles lacks application. For the argument of essay 3 has the consequence that the state conceptualist himself is committed to there being such representational capacities. Indeed, it has the consequence that he is committed to there being *conceptual* representational capacities of this kind. Of course, we have not argued that it would be

right to use the Kantian or Systematicity Principle to do this job. But the possibility of this type of reply gives grounds for optimism that a good reply of broadly the same kind will be available to us to address the question how perceptual experience can be contentful yet not conceptualised.

Even if one grants that state nonconceptualism is compatible with representationalism, one might doubt its compatibility with the stronger doctrine of content conceptualism. One might think, that is to say, that state entails content nonconceptualism (contrapositively, that content entails state conceptualism). If this is so, it bears on essay 4. That essay concludes that content conceptualism is incompatible with the thesis of diaphanousness, (D), without assigning blame to either. If content conceptualism is false for independent reasons, it would lessen the dialectical pressure on (D).

However, it is not obvious that state entails content nonconceptualism. The issue critically depends on what kinds of content beliefs and thoughts have. If the contents of beliefs are of a moderately or extremely coarse-grained sort – e.g. take form of neo-Russellian propositions (Salmon 1986; Soames 1987) or of functions from possible worlds to truth-values (Stalnaker 1984, 1999) – it is not at all clear why perceptual experiences and beliefs could not have contents of the same sort even though they differ in point of conceptualisation. Suppose, for example, that Tye's (2000, 2006a, 2006b) view of experience is correct, on which perceptual experiences have neo-Russellian propositions as content, and involve the activation of certain nonconceptual representational capacities.³⁶ If beliefs also have neo-Russellian propositions as

³⁶ For Tye, the relevant perceptual-representational capacities seems to count as nonconceptual for the reason that they do not connect with, nor place demands on, our memory systems in the way concept-possession does.

contents, the content of an experience is the content of a possible belief provided, roughly, that one can attribute in a belief just the properties one experiences things as having to just the objects one experiences as having them. Idealising away from performance limitations (in a way that seems entirely legitimate in this context), there seems to be every reason to think one can have beliefs of this sort.³⁷ The assumption that beliefs involve conceptual representational capacities in no way disturbs this verdict. If the contents of beliefs are of a neo-Russellian sort, then, there seems to be no reason to think state nonconceptualism should rule out content conceptualism. The same conclusion seems if anything better founded if the contents of beliefs take an even more coarse-grained form, such as functions from possible worlds to truth-values.

What if the contents of beliefs are of a more fine-grained kind, such as neo-Fregean Thoughts, i.e. structured complexes of modes of presentation of objects and properties, individuated in terms of considerations of cognitive significance?³⁸ If this is what the contents of belief are like, content conceptualism would imply that perceptual experiences have contents built up from modes of presentation of exactly the same type. One might think however that modes of presentation are typed in a way that track relevant differences in representational capacities. If perceptual experiences and beliefs involve representational capacities so different as to make the latter but not the former qualify as conceptual, one might therefore conclude that this difference would correlate with a difference of type in mode of presentation. That would block content conceptualism. So, one might think, on this neo-Fregean view of the content of beliefs, there is no way of combining content conceptualism with state nonconceptualism.

³⁷ See footnote 2, essay 4 for a comment on the need for the relevant idealisation, and its justification.

³⁸ For such neo-Fregean views, see among others Evans (1982), Peacocke (1992) and McDowell (1998c).

I grant that the case for an entailment from state to content nonconceptualism is stronger on a fine-grained, neo-Fregean view of content. However, it all hinges on the finer details of how the things belonging to the critical category of a mode of presentation are individuated. If their individuation is highly unconstrained, the implication from content to state conceptualism can trivially be secured by ruling that ‘mode of presentation of the sort entertained by conceptual representational capacities’ and ‘mode of presentation of the sort entertained by nonconceptual representational capacities’ denote two mutually exclusive types. Once tighter criteria of individuation are laid down, the implication becomes non-trivial. Suppose, for example, that one adopts as criterion what is often referred to as ‘Frege’s Principle’, i.e. the principle that content ...*F*... and content ...*G*... (obtainable from ...*F*... by substituting *G* for *F* at one or more places) are distinct if it is possible for someone rationally to believe ...*F*... while doubting or denying ...*G*.... Suppose, moreover, that one adopts a principle according to which it is characteristic of conceptual (as opposed to nonconceptual) representational capacities to presuppose the understanding a word in a public language that expresses the content associated with the representational capacity. On these suppositions, it is far from clear why perceptual experiences could not have representational content of a type entertainable in belief even if the beliefs differ from experiences in that the former but not the latter involve representational capacities linked with expressibility. Of course, the suggested distinguishing mark on conceptual representational capacities may not be very plausible. That said, it will require substantial work to show that content conceptualism entails state conceptualism on a neo-Fregean conception of content. Of course, it is an equally substantial task to show that content does *not* entail state conceptualism on such a conception of content. In my knowledge of the literature, such

a project has not been undertaken, and pursuing it here would take us much too far afield. Antecedent to such a project, the most justified stance would seem to be to leave as an open question whether state entails content nonconceptualism.

6.3 The bearing of the incompatibility of content conceptualism and (D)

Essay 4 argues that there is an incompatibility between content conceptualism and the thesis of diaphanousness, (D), without saying which is false (of course both may be). I will now consider what bearing the falsity of either of these has on the other essays in the dissertation.

Let's suppose that content conceptualism is false, i.e. that content nonconceptualism is true. This would be bad news for the argument of essay 1 if it turned out that content nonconceptualism was incompatible with representationalism. Do we have any reason to be worried about such incompatibility? Over the last twenty-odd years a number of philosophers, notably Peacocke (1989, 1992, 2001) and Tye (2000, 2006a, 2006b), have developed views on which perceptual experiences have representational content even though the representational contents they have are of a different type than the content of thoughts and beliefs, subject to a non-equivalent principle of individuation. For example, on Tye's view, the contents of experience have the more coarse-grained individuation of neo-Russellian propositions while the contents of beliefs have the more fine-grained individuation of neo-Fregean Thoughts. It seems hard to deny that both of these notions are coherent notions of a type of representational content: they are notions of something that determines truth-condition and captures facts about how a given state represents the world as being. Of course, Peacocke's and Tye's proposals may be wrong on many points of detail. And, of course, their proposals may be wrong for the reason that content conceptualism is true – not false as we are currently supposing. My point

here is only that the very idea of combining content nonconceptualism with representationalism does not seem to be incoherent. Absent further arguments to the contrary (which to my knowledge have not been given in the literature), we are entitled not to be too worried about content nonconceptualism being incompatible with representationalism.

Even if content nonconceptualism is compatible with representationalism one might think it isn't compatible with state conceptualism. One might think, that is to say, that content entails state nonconceptualism (contrapositively, that state entails content conceptualism). This is, of course, the converse entailment of that considered in the previous subsection. If this entailment held good, it would seem to be, if anything, good news to the argument of essays 2 and 3. The drift of these essays is strongly towards state nonconceptualism. If content nonconceptualism is true, and the entailment from content to state nonconceptualism holds, we would have independent support for that conclusion.

We have already observed one reason why state does not entail content conceptualism. It is because the latter entails representationalism (where representational content is understood as propositional, i.e. truth-conditional, content) whereas the former neither entails nor presupposes representationalism (understood in that way), as we saw in section 6.1. The question we need to ask is rather whether state conceptualism, in conjunction with representationalism, entails content conceptualism. Now, we take representationalism to be true and independently justified. If content nonconceptualism in conjunction with representationalism entails state nonconceptualism, we would thus still get the 'friendly implication' in favour of state nonconceptualism.

Just as we found when we asked whether content entails state conceptualism, the question whether state conceptualism cum representationalism entails content conceptualism hinges on issues that are delicate and controversial. Notably, it hinges on the question by appeal to what principles one distinguishes conceptual from nonconceptual representational capacities, and how these principles constrain the kinds of content that can be associated with these capacities. Going into these issues would take us too far afield. Nevertheless, it would seem natural to think that if perceptual experience, just as much as beliefs, gets its content through the exercise of conceptual capacities (and that content is propositional content), then surely one could entertain that very same content in belief too. At the very least, it would seem that the burden of proof is on the theorist who affirms that this content conceptualist conclusion does not follow. There may be some grounds for thinking, then, that the ‘friendly implication’ holds.

Content nonconceptualism also connects with the discussion of Brentano’s thesis in essay 6. In particular, some influential versions of content nonconceptualism connect with the judgement made at the end of that essay, that it is doubtful whether simple perceptual experiences meet the strong necessary conditions for intentionality to which the defender of Brentano’s thesis would seem obliged to resort. According to one of these strong conditions, a mental state is intentional only if it is ascribed by what we call a ‘trans-Russellian’ report. This is a report where the content-specifying terms figure in a context that is susceptible to substitution failure even for expressions having (in a sense there defined) the same ‘Russellian meaning’ (roughly, the same semantic

structure and the same referents for the atomic constituents).³⁹ Beliefs, desires, and other paradigmatic propositional attitudes may seem to meet this necessary condition, insofar as it appears

(5) Lois Lane believes Superman is strong.

and

(6) Lois Lane believes Clark Kent is strong.

may differ in truth-value.⁴⁰ However, as I suggest in essay 6, it is less clear whether bodily sensations and visual experiences meet this condition. I there consider a case involving reports of bodily sensations. I want for the nonce however to look at a case involving reports of visual experiences – a case that once again may indicate a contrast with the case of belief reports. Suppose Frank is sitting on a beach on Rhodes, gazing out at the Aegean. Frank is in a position to refer demonstratively to the colour the water looks to him. In particular, it is true that

(7) Frank believes that the water is *that* colour.

and that

(8) The water looks *that* colour to Frank.

In interpreting the demonstrative ‘*that* colour’ in (7) and (8), we are to pretend we are occupying a visual perspective relevantly like Frank’s, and so are in a position to make the demonstrative reference he makes. Now, the colour the water looks to Frank, and to which his demonstrative refers, is turquoise. Yet Frank is under the misapprehension that turquoise is a shade of violet. Under these circumstances, it may seem that

(9) Frank believes that the water is turquoise.

³⁹ See section 5, essay 6, for more on this notion.

⁴⁰ I do not want to endorse the claim that they may. The situation is complicated, since the putative appropriateness of (3) as a report of the belief Jimmy expresses by saying (1) suggests that co-referential terms may be substituted in belief contexts. Again, see Richard (1990) for a rich discussion of the various complexities surrounding belief reports.

is false. For Frank (let's suppose) believes that the water is not a shade of violet. Thus if he did believe that the water is turquoise, then given his belief that turquoise is a shade of violet and his basic logical competence, he would also have believed that the water is a shade of violet. In other words, if he did believe that the water is turquoise, we would have had apparently obviously contradictory beliefs. Yet the simple misapprehension Frank is under does not make him irrational. So it would seem that (9) is false. In contrast, many have the intuition that

(10) The water looks turquoise to Frank.

is true, in spite of his misapprehension. This may suggest that reports of visual experiences, in the shape of what we call 'non-comparative qualitative looks reports' in essay 1, allow for substitution of terms which refer to the same property, and thus that they are not trans-Russellian, in contrast with belief reports. Tye (2002, 2003) takes this putative contrast to be a reason in favour of his content nonconceptualist view that the content of visual experience is subject to more coarse-grained individuation conditions than that of beliefs. I do not want to take a stand here on whether Tye's argument is convincing.⁴¹ I note it here as an illustration of an issue where the questions of content nonconceptualism and Brentano's thesis become closely linked.

Let's suppose, on the other hand, that (D) is false. One drastic suggestion would be that this is indicative of widespread error in the considerations of diaphanousness with which (D) is associated as but one salient claim. This would be a dramatic consequence indeed, as a great number of influential arguments in contemporary philosophy of consciousness are based in one way or another on considerations of diaphanousness,

⁴¹ For more on this argument of Tye's, see section 7.2 below.

notably arguments against various types of spectrum inversion scenarios.⁴² However, the arguments in essay 4 do not give good reason to suspect widespread or pervasive error in the considerations of diaphanousness, even if we presume that content conceptualism is correct. To see this, consider the following claim:

(D') In introspection, one is or becomes aware of the phenomenal character of visual experience only *by* attending to, or being aware of, the fact that things look such-and-such ways to one.

(D') differs from (D) in that the target of the attention (or awareness) by means of which one becomes aware of phenomenal character is not the *objects* that appear to one and the *properties* they appear to have but the *fact that* things looks such-and-such ways to one.⁴³ It still qualifies as a claim of diaphanousness in a generous sense, in that it reasonably can be construed as excluding that one becomes aware of the character of the experience by attending to or being aware of the fact that experience has such-and-such intrinsic features. At any rate, it excludes this if for the experience to have such-and-such intrinsic features is something else than for it to be one in which things look such-and-such ways to one. Yet, even if we suppose content conceptualism is correct, (D') does not give rise to the problem (D) was found to give rise to. When I come to have a visual experience at *t* (having had no relevant visual experience before *t*), then even if I had had a thought with the same representational content up until *t*, there is at *t* the change that it comes to be so that things looks such-and-such ways to me. It is perfectly

⁴² See Thau (2002) and Shoemaker (2003) for but two examples.

⁴³ It also differs from (D) in allowing that the mode of access to the relevant target, by means of which one is aware of phenomenal character, takes the form of awareness rather than attention. This is to appease those readers who may harbour doubts about the notion of attention to facts. I don't believe such doubts are very serious, but in any case formulate (D') in a way that it neutral on the matter.

consistent with (D') that I come to be aware of this fact. If I become aware of the fact that things look such-and-such a way to me, it is reasonable to conclude that I thereby become aware of my state of mind as having a visual or perceptual phenomenal character. As I argue in section 2, essay 4, if I am aware of my state of mind as having such a visual or perceptual phenomenal character, then we should certainly expect me to be able to tell that there was a change in the phenomenal character of my state of mind when I shifted from merely having the thought to (perhaps also) having the experience. The problem we identify in section 2, essay 4, is that it is hard to see how one could come to be aware that one's state of mind has a perceptual phenomenal character if (D) is correct. This is not a problem on (D'), and thus we have no reason to think (D') incompatible with content conceptualism.⁴⁴

I cannot see that any of the claims or arguments of the other essays in the dissertation presume that (D) as opposed to (D') is true (to the extent that they indeed are sensitive to the truth of either). If (D) is false, and false just for the reason that content conceptualism is true and the arguments of essay 4 sound, it would thus not seem to have an important bearing on the other essays in the dissertation.

7. The relation of this work to *Content in Thought and Perception*

The essays collected here in part complement, in part reinforce, and in part raise critical questions about the argument of my 2006 thesis *Content in Thought and Perception* (henceforth *Content*). In this section, after summarising that thesis, I trace its

⁴⁴ See Siewert (2004) for a recent interesting discussion of diaphanousness, congenial to the idea that something like (D') is a more promising claim of diaphanousness than something like (D). (Siewert's paper came to my attention after essay 4 had been written.)

connections with the essays collected here, emphasising issues of particular critical import to the argument of *Content*.

7.1 Précis of Content

The principal thesis defended in *Content* is a claim I label ‘content congruence’ and formulate as follows:

Content congruence The content of a perceptual experience is the content of a thought a reflective subject can think if she has the experience.

This formulation is intended to encapsulate two claims. First, the claim that any perceptual experience of the sort had by reflective subjects (i.e. subjects capable of self-conscious thought, such as normal human adults) have a representational content, where the content in question is identical to that of a possible belief or thought. Content congruence was thus intended to entail content conceptualism, restricted to the kinds of perceptual experiences reflective subjects have.⁴⁵ It was restricted in this way because I wanted to remain neutral on whether the perceptual experiences of animals not capable of self-conscious thought are of the same kind, as far as their contents are concerned, as the perceptual experiences of self-conscious creatures. We may disregard this restriction here.

⁴⁵ As I note in fn. 4, essay 4, the claim of content conceptualism is susceptible of a stronger and weaker reading. The stronger reading says that each perceptual experience *exclusively* has a representational content such that it is the content of a possible belief or thought. In other words, it does not also have a representational content that is *not* the content of a possible belief or thought. The weaker reading says the same apart from the ‘exclusively’ rider and is thus consistent with the claim that some or any perceptual experience also has representational content which is not the content of a possible belief or thought. Content congruence was intended to entail content conceptualism on the stronger reading. However, in retrospect, I believe some of the arguments presented at the time as they stand succeed at most in sustaining content conceptualism on the weaker reading. I return to this. (Of course, if perceptual experiences have representational content of but one kind, the strong and weak readings are equivalent.)

The second claim I intended content congruence to entail presupposes but goes beyond content conceptualism. Take an arbitrary perceptual experience and a possible thought with the same content as it. This second claim says that a reflective subject who has the former will be in a position to have the latter. On the intended reading, it is consistent with this claim that it may not be psychologically possible for the subject who has the experience to have the corresponding thought if she suffers from various performance limitations, e.g. if she is currently preoccupied with other tasks, or scared witless, or lacking in motivation, etc. etc. I believe this second claim is plausible if content conceptualism is, but I will not emphasise this point, and concentrate on the issue of content conceptualism (as I did in *Content*).

Content gives two lines of argument for content conceptualism. The first, ‘simple’ argument (ch.1) relies on the claim that each perceptual experience is state in which it perceptually appears to the subject that things are thus and so; the content of the experience is precisely what perceptually appears to the subject to be the case in it, viz., that things are thus and so. Corresponding to such a perceptual experience there is a possible thought, viz., the thought that things are thus and so. The content of this thought is what is thought in thinking it, viz., that things are thus and so. So the content of the thought is the same as that of the perceptual experience.⁴⁶

⁴⁶ I conceived of this ‘simple’ argument as a way of articulating a line of reasoning suggested by such remarks as the following from McDowell: ‘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 that already has conceptual content. In experience one takes in, for instance sees, *that things are thus and so*. That is the sort of thing one can also, for instance, judge.’ (1994: 9, emphasis in original)

The second argument (ch. 4) for content conceptualism appeals to the role of perceptual experience in psychological explanation of observational beliefs. I argue that this explanatory role presumes a non-vacuous *ceteris paribus* psychological law linking having a perceptual experience with a given content with the acquisition of a belief with a correlated content. I contend further that the content conceptualist is significantly better placed to formulate such a law, consistently with his view, than his content nonconceptualist opponent. If content conceptualism is correct, the following is plausible as at least a first approximation to such a law:

The Conceptualist Law *Ceteris paribus*, for any reflective subject *S*, for any content *p*, if *S* has an experience with the content *p*, *S* is interested in whether *p*, *S* does not believe that there is anything abnormal about the conditions of perception, then *S* believes *p*.

Now, if content nonconceptualism is right, this cannot be a non-vacuously true law.⁴⁷ For it follows from content nonconceptualism that if a content is the content of an experience, it cannot also be the content of a belief. From the point of view of content nonconceptualism, it seems a formulation of a law of the relevant must take rather the following form.

The Nonconceptualist Law *Ceteris paribus*, for any reflective subject *S*, for any perceptual content *p*, for any thought content *q*, if *S* has an experience with the

⁴⁷ Here, content *nonconceptualism* must be taken as the denial of content conceptualism on the weaker reading (see footnote 45 above). Content nonconceptualism must, in other words, be taken to be the thesis that perceptual experiences exclusively have representational contents of a kind such that they are not the contents of possible beliefs or thoughts. Correlatively, the present argument supports, at best, content conceptualism on the weaker reading.

content p , q is a conceptualisation of p , S is interested in whether q , S does not believe that there is anything abnormal about the conditions of perception, then S believes q .

The crucial difference between this and the Conceptualist Law is the use it makes of the notion of a thought content being a ‘conceptualisation’ of a perceptual content. The question is, then, how this relation of conceptualisation is to be understood. I consider a number of the answers a content nonconceptualist would be in a position to offer, but argue that they either render the Nonconceptualist Law false or suffer from either vicious circularity or dialectically illegitimate appeal to the *ceteris paribus* clause. The difficulty of this challenge of giving an account of the relation of conceptualisation is a reason, I maintain, to prefer the content conceptualist view.

A major part of *Content* is devoted to responses to four important criticisms of content conceptualism. One of these is due to Peacocke (1986, 1989), and argues that whereas the content of thought is individuated in accordance with ‘Frege’s Principle’ (see sections 3 and 6.2 above), the content of perceptual experience is not. This argument, I contend (ch. 3), fails to establish content nonconceptualism since the version of Frege’s Principle that perceptual content fails to conform to is one to which the content of beliefs and thoughts also fails to conform. Another objection is also due to Peacocke (1992). It makes the case that the conceptualists will not be able to provide an adequate, non-circular, individuating account of observational concepts in terms of the canonical role they play in the cognitive life of thinkers possessing them. Several conceptualists, including McDowell (1994), Sedivy (1996) and Brewer (1999), have rejected this objection on the ground that Peacocke’s non-circularity requirement is illegitimate and

unjustified. I argue (ch. 5) that there is an alternative reply available to the conceptualist, on which there is no need to make such a contentious claim and all of Peacocke's central constraints on concepts can be respected. This alternative reply says that an account of the sort Peacocke demands can be given if we appeal, among other things, to (i) causal links between perceptual experiences involving a given concept and features of the environment that gives rise to such experiences, and (ii) links between thoughts involving a given concept and dispositions to action.

A third line of objection to content conceptualism, and indeed to representationalism more generally, is developed in Campbell (2002). His argument relies on the idea that perceptual experience of our surroundings has a crucial explanatory role to play in an account of how we achieve knowledge of reference of demonstratives referring to things around us. Having such knowledge of reference is essential since it is what causes and justifies the use we make of demonstratives in our basic empirical thought and talk about the world. The conceptualist will not be able to honour this explanatory role – indeed that no representationalist will – since he essentially takes for granted what is to be explained, viz., the intentionality of basic empirical cognition. I distinguish several sub-arguments in Campbell's case here, and argue that the conceptualist is in a position to meet them (ch. 6).

A fourth objection, implicit in various remarks of Tye's (2000, 2002, 2003), is of special relevance to at least one of the essays collected here. The objection starts from the premise that, in order to get at the nature of visual experiences, we should rely not on reports of the form 'It visually appears to S that P', as maintained by the 'simple' argument outlined above, but on reports of the form 'X looks F to S', where 'F' holds

the place for an expression for a 'sensible property', i.e., a property 'of which one is directly aware via introspection as one undergoes a sensory experience' (Tye 2000: 54). In reports of the latter type, 'F' has the job of specifying the representational content of the visual experience so reported. This means that the question of what kinds of content visual experiences have can be approached by looking at the semantic functioning of 'F' in this construction, just as the same question for beliefs can be approached by looking at the semantic functioning of 'F', or of 'a', in reports of the form 'S believes that a is F'. When we do so, we find an asymmetry between the visual and the doxastic case, as we suggested was illustrated by the case of Frank, looking out on the Aegean, in section 6.3 above. This contrast is indicative of the fact that visual experiences have contents with the more coarse-grained conditions of individuation of neo-Russellian propositions while beliefs have contents of the more fine-grained conditions of individuation. Or so Tye argues.

My response to this objection in *Content* (ch. 2) is to take issue with Tye's contention of a relevant contrast, in point of substitution resistance, between belief reports and looks reports. A major – perhaps crucial – reason for denying substitutivity in belief and desire reports is the idea that affirming substitutivity will lead to explanatory loss in common-sense, rational explanation of the actions and attitudes of people in terms of their reasons. For example, Lois Lane's sudden urge to get to Union Square by noon is better explained by her desire to catch a glimpse of Superman, and her belief that Superman will be there at noon, than it is by her desire to catch a glimpse of Superman and her belief that Clark Kent will be there at noon. This sort of reason for denying substitutivity, I suggest, generalises to the relationship between visual experiences, as expressed by looks reports, and beliefs. Let's suppose a type-physicalist view of colours

is correct, according to which colours are identical to types of surface spectral reflective profiles.⁴⁸ Let's also assume green in particular is identical with a surface spectral reflectance type for which the technical term 'optizygry's catches on in the vision science or optical science communities. My claim, then, is that Tom's visually acquired belief that the paper is green is better explained by the paper's looking green to him and his taking the experience at face value, than it is by the paper's looking optizygry's to him and his taking the experience at face value. At any rate, my claim is that there is a putative explanatory contrast here that corresponds to the putative explanatory contrast noted in the Lois Lane case. Thus, if the latter gives good reason to deny substitutivity for 'F' in the context 'S believes that ...F...', the former gives good reason to deny substitutivity for 'F' in the context 'X looks F to S'. I allow that the reason may not be conclusive in either case. No attempt is made to refute the responses such neo-Russellian theorists as Soames (1990) and Braun (2000) have offered to the argument from the explanatory role of psychological reports to the denial of substitutivity. My claim is only that if these sorts of responses work for the 'looks' case, they also work for the 'believes' and 'desires' case. The contrast Tye relies upon thus does not hold. This, then, was my response to Tye as set out in *Content*. I shall come back to this in a moment.

The final chapter of *Content* (ch. 7) added an objection to state conceptualism to the defence of content conceptualism in the earlier chapters. I there formulated the claim under attack as follows:

⁴⁸ See Byrne & Hilbert (1997) for a defence of such a view.

State conceptualism (*Content*) The concrete representational capacities in virtue of the activation of which a perceptual experience has the content it has are conceptual.⁴⁹

This formulation obviously diverges from our canonical formulation of state conceptualism here, viz.,

State conceptualism (Essays 2 and 3) In order to have a perceptual experience as of ...F... a subject needs to possess a concept of F/F-ness/F-hood...

To keep these two formulations apart, let's label them the 'Constitutional' and the 'Requisitory' formulations, respectively.⁵⁰ The Constitutional differs from the Requisitory in being tacit on how the specification of the content of the experience correlates with the identity of the representational capacities involved. For all the Constitutional says, a perceptual experience as of something being triangular may rest on the activation of a conceptual representational capacity to represent something as porridge. However, I take a link of the sort incorporated by the Requisitory to be independently plausible. Given this presupposition, the Constitutional plausibly entails the Requisitory. If a perceptual experience is a perceptual experience as of ...F..., rather than an experience as of, say, ...G..., thanks to the activation of a conceptual representational capacity to represent F, then having a perceptual experience as of

⁴⁹ The distinction between 'concrete' and 'abstract' representational capacities is explained in ch. 7 of *Content* but need not concern us here.

⁵⁰ The constitutional formulation was inspired, among other things, by such remarks as the following from McDowell: '[W]e need a conception of experiences as states or occurrences that are passive but reflect conceptual capacities, capacities that belong to spontaneity, in operation' (1994: 23). If experiences 'reflect conceptual capacities ... in operation' it is presumably because they have the content they have in virtue of the activation (operation) of these capacities.

...F... requires a conceptual representational capacity to represent F, i.e. requires possessing a concept of F. However, the converse entailment does not hold. One could hold a view consistent with the Requisitory on which the fact that a perceptual experience has a certain content requires, and is explained by, the possession or activation on the part of the subject of *both* nonconceptual and conceptual representational capacities. So the Constitutional is stronger than the Requisitory (contrapositively, Requisitory state *nonconceptualism* is stronger than Constitutional state *nonconceptualism*).

The argument against Constitutional state conceptualism in *Content* builds on Heck's (2000) objection against conceptualism. Heck considers what account a conceptualist can offer of the fine-grained content of experience in a case of illusion – in a case where things are not as they appear. He argues that the available options either falsely imply that the experience is veridical (since they invoke a perceptual-demonstrative, conceptual capacity that refers to a property the object seen in fact has) or suffer from vicious circularity (since they presume a priori *nonconceptual* representational capacity, activated in the experience, on which the conceptual capacity invoked rests). Heck's argument has, however, been interestingly criticised by Pelling, who argues that the conceptualist is, in the end, no worse placed than his nonconceptualist opponent to give an account of illusory experience.⁵¹ I seek to add strength to Heck's line of argument in the face of these criticisms, appealing to some general constraints on reference determination that the nonconceptualist will be better placed to meet than his conceptualist opponent.

⁵¹ Pelling (under review), based on his presentation, Pelling (2005).

The final sections of that chapter were concerned with whether the state nonconceptualism there defended is consistent with the content conceptualism defended in the earlier chapters. I point out, as I did in section 6.2 above, that there is no reason to expect an entailment from content to state conceptualism on a moderately or extremely coarse-grained conception of the content of beliefs. I also argue (in a way that overlaps with but expands upon that in section 6.2) that the entailment is far from obvious even on a more fine-grained, neo-Fregean conception of the content of belief. The sorts of considerations that underlie the judgement that the representational capacities operative in experience are nonconceptual does not, or at least not clearly, imply that the cognitive significance of perceptual content differs from the cognitive significance of the content of perceptual judgement immediately based on experience.

7.2 Relations

In *Content*, I operate throughout on the assumption of representationalism, or what I there call ‘weak intentionalism’. While I devote three or four pages (in section 1.2) to motivate this assumption in broad outline, it effectively figures as a presupposition, something taken for granted. Essay 1 here explores the case for representationalism in much greater detail, by way of responding to Travis’s objections to it, which were not addressed at all in *Content*. It thus effectively carries out groundwork for the project of *Content*.

Chapter 7 of *Content* gives an argument for state nonconceptualism, appealing to highly general constraints on reference determination. Essays 2 and 3 in effect reinforce this

conclusion, in as much as they give additional, independent arguments the drift of which is friendly to that view.⁵²

Essays 4, 5 and 6 address problems and issues that *Content* does not go into, although they are certainly relevant to it. Essays 4 and 6 in particular at certain points raises critical questions about the argument of *Content*. Essay 4 does so by arguing that content conceptualism, supported in *Content*, is incompatible with the thesis of diaphanousness, (D). Now, in essay 4, as in this Introduction, I am officially neutral on whether this incompatibility reflects falsity on the part of (D), content conceptualism, or both. Were I to abandon my perch on the fence, however, my verdict would be that (D) is false. I am inclined towards this verdict as I am inclined to think that the style of argument for incompatibility between (D) and content conceptualism developed in essay 4 can be generalised to show incompatibility between (D) and content *non*conceptualism. To explore this generalisation is a project I aim to pursue in the near future.⁵³

⁵² To be precise, essays 2 and 3 point out problems for what we called Requisitory state conceptualism, and to that extent are supportive of the denial of that view, Requisitory state nonconceptualism. Since Requisitory state nonconceptualism entails Constitutional state nonconceptualism, they are to the same extent supportive of the latter.

⁵³ The following briefly suggests the sort of generalisation I have in mind. Even if content nonconceptualism holds, it is plausible that one can think about the same objects as one experiences and, in thinking about them, attribute to them the properties they appear to have in the experience. At the level of the objects upon which the mind is directed, and the properties attributed to them, the experience and the thought thus agree. If we compare a Wall-Experience and a Wall-Thought related in this manner, it seems that the problem of explaining in terms of (D) how the phenomenal difference between them will be introspectively detectable remains in force. For according to (D), one gets at the character of the experience only by attending to the *objects* that appear and the *properties* they appear to have. The fact that they *appear* such-and-such is not among the things by which one, in attending to them, gets at the character of experience.

Essay 6 connects with the argument of *Content* through a remark made at the end of the essay. I there ask whether standard reports of visual experiences and pains are ‘trans-Russellian’, i.e. whether they are susceptible of substitution failure even for names referring to the same object or for predicates referring to the same property. I write that, even assuming that paradigmatic attitude reports are trans-Russellian, there is an intuition that pain reports allow for substitution of terms referring to the same object or property. Indeed, I say this intuition ‘is certainly not without force’. If these remarks are apposite about pain reports, the same remarks are plausibly apposite about such visual reports as ‘X looks F to S’. Now, as we saw, in chapter 3 of *Content*, I argued, *pace* Tye, that there is no relevant contrast in point of substitution resistance for ‘F’ between ‘X looks F to S’ and ‘S believes that ...F...’. So the remarks at the end of essay 6 may seem to be in direct conflict with an important claim in *Content*.

Now, there is no direct contradiction here, for two reasons. First, neither in essay 6 nor in the discussion of its relation to content conceptualism above do I commit myself on the trans-Russellian nature or otherwise of paradigmatic attitude reports. If they are not trans-Russellian, then the admitted forcefulness of the intuition that looks reports are non-trans-Russellian is not even a *prima facie* problem for the defence of content conceptualism in *Content*. Second, even if paradigmatic attitude reports are trans-Russellian, I shy away from endorsing outright that looks or pain reports are non-trans-Russellian, admitting only that there is a forceful intuition to this effect.

Going into the delicate question of whether to endorse the forceful intuition – or not, should the argument of chapter 3 of *Content* prove otherwise – would take us too far afield here. Nevertheless we may at least say the following. There is a link between the

fate of the criticism of Brentano's thesis in essay 6 below and the defence of content conceptualism in chapter 3 of *Content*. If attitude reports are trans-Russellian and looks reports are too, then so much the better for the argument of *Content*, chapter 3; at the same time, there is one less problem for a defence of Brentano's thesis than essay 6 suggests. On the other hand, if there is the semantic contrast between looks reports and belief reports Tye claims there to be, while this is bad news for the argument of *Content* it still confirms one of the problems suggested for a defence of Brentano's thesis in essay 6.

Essays 2 and 3 are also raising a critical question of the argument of *Content* – to wit, whether the state nonconceptualism they (the essays) defend (or at least indirectly support) is compatible with the content conceptualism defended in *Content*. As already noted, this critical question arises with equal force internally to *Content*, as chapter 7 of the latter defends state nonconceptualism. I have indicated above why I hold there to be no evident reason to think that content conceptualism is incompatible with state nonconceptualism. However, by my own admission in *Content*, I do not there show that these claims *are* compatible. This was one of the admitted weaknesses of *Content*, which I hoped to rectify in future work. Unfortunately, as indicated in section 6.2 above, the opportunity to take on this task has yet to present itself. This is not a great problem for the essays collected here, as they may all be read in a spirit of official neutrality on the question of content conceptualism. All the same, it would of course be nice for the overall coherence of these essays with *Content*, and *Content* with itself, if such compatibility could be shown. I can do no more here than express again a hope that circumstances will allow me to pursue this task in the not too distant future.

8. Concluding remark on contribution

The essays collected in this dissertation engage with current debates in the philosophy of perception and philosophy of mind more generally. They address in particular discussions of the connections between experience and intentionality, experience and concepts, intentionality and phenomenology, and intentionality and mentality. Their contribution to these debates in a nutshell is

- to enhance our understanding of, and improve the case for, the thesis that perceptual experiences have representational content, by way of responding in detail to Travis's important criticism of the thesis;
- to improve the case against McDowell's suggestion that subjects have recognitional capacities for the fine-grained colours, shapes or other features they experience things as having;
- to clarify the commitments of (state) conceptualism by showing that these include the rejection of two widely granted principles about concepts – the Kantian and the Systematicity Principles – and thereby throw light on the question what the principled distinction between conceptual and nonconceptual representational capacities may be;
- to identify a problem about the compatibility of content conceptualism with a certain much-discussed and fairly popular thesis of diaphanousness, (D);

- to outline an object-directed counterpart to Jackson's experience-directed knowledge argument, and show how the close parallels between these mean that anyone who endorses a 'perspectivalist' account of the latter ought to endorse a corresponding 'perspectivalist' account of the former, where this has the consequence that close analogues of introspective-phenomenal concepts figure in basic perceptual judgement;
- to ask whether intentionality is sufficient to mentality, and show how reflection of this question leads to a significant, underappreciated challenge to a defence of Brentano's thesis.

Looks and Representational Content: A Defence of the Orthodoxy

The orthodox view in contemporary philosophy of perception is that perceptual experiences have representational content.¹ The orthodox disagree among themselves as to what form the representational contents of experience take, what features of the world they are capable of representing, and what features of perceptual experiences they are capable of accounting for.² They agree, however, at least on the following. If perceptual experiences have representational content, they are analogous to beliefs and desires in at least the following respect. Beliefs have truth-conditions, and desires fulfilment-conditions. There is a certain notion of *content*, let's call it representational content, applicable to beliefs and desires, that accounts for this fact: the truth-conditions of a belief, or the fulfilment-conditions of a desire, are determined by the representational content of the belief, or the desire. Analogously, visual experiences have veridicality-conditions. There is a notion of content, viz. representational content, applicable to visual experience that accounts for this fact: the veridicality-conditions of an experience are determined by its representational content. Chalmers (2004) refers to

¹ See Chalmers (2004) and Siegel (2005) for two excellent surveys of the orthodox view, in its myriad varieties.

² For opposing views on the first issue, see e.g. Peacocke (1992) and Brewer (1999); on the second issue, e.g., Siegel (2006) and Clark (2000); on the third, e.g., Peacocke (1983) and Tye (1992).

this orthodox view as ‘weak representationalism’. I will follow him in this, aside from dropping ‘weak’ for brevity.³

In light of the by now orthodox status of representationalism, it is unsurprising, perhaps, that it has of late increasingly attracted dissenters. In this paper, I will examine Charles Travis’s (2004) recent frontal assault on the view.⁴ Travis’s criticism is nothing if not fundamental. He argues that there is ‘nothing in perceptual experience which makes it count as having some one representational content as opposed to countless others’ (p. 85) and that ‘[p]erceptual experience is not as such either veridical or delusive’(p. 57).⁵ In the response to follow, I discuss in detail what I take to be Travis’s two principal lines of argument against representationalism, and make the case that they fail. I should note that although I take these two lines of argument to be the burden of Travis’s case,

³ Interestingly, representationalism has not been orthodox for very long. As late as around the mid-century, if not later, the most popular philosophical accounts of sense experience were the sense-datum theory (the default view, as it were) and the adverbial theory (its perceived main rival). It should be noted, at this point, that many sense-datum theorists in effect accepted, with qualification, the defining claim of representationalism, since they held that the presentation of a sense-datum inevitably was accompanied by some sort of interpretation, by which the sense-datum was taken to signify something in the physical world. This accompanying interpretation clearly had representational content on their view. However, they denied that the ‘sensory core’ of sense experience, consisting in the presentation of sense-data, had any representational content. The same structure of ‘sensory core’ vs. ‘interpretation’ would typically be incorporated on adverbial theories. Representationalists differ from these theorists either in denying that there is any relevant distinction between sensory core and interpretation within perceptual experience, or maintaining that even the sensory core has representational content. See Firth (1949-50) for an early systematic discussion of, and criticism of, the distinction between sensory core and interpretation.

⁴ Other recent critics of (different aspects of) representationalism include Campbell (2002), Brewer (2006, 2007, forthcoming) and Breckenridge (2007). Johnston (2004) makes critical remarks about the idea that experience is ‘propositional’, and may seem to distance himself from what he calls ‘Intentionalist accounts’. However, as we shall see below, his view is not really an alternative to representationalism in our sense. For some objections to Brewer’s anti-representationalist account of visual illusion, see nt. 41 below. I respond to Campbell’s objections to representationalism in Nes (2006: ch. 6).

⁵ All un-attributed references are to Travis (2004).

there are however various subplots in his extraordinarily rich and complex paper I will not be able to address.⁶

The first of the principal lines of argument, to which I devote by far the most attention, trades on the claim that facts about how things look to a subject cannot ground the ascription of representational content to his visual experience. Within this first line of argument I shall, in turn, distinguish two formulations. The formulations are closely related, and aim in each case to show that facts about how things look determine an incoherent content for perceptual experiences if they determine any. However, the formulations raise subtly different issues that benefit from separate treatment. I tackle the first formulation in sections 3 and 4, and the second formulation in section 5. My overall conclusion is that Travis's arguments leave unharmed the idea of there being a class of non-comparative, qualitative facts about the way things look, such as the fact that something looks blue to one, that ground the non-arbitrary ascription of representational content.

In section 6, I discuss Travis's second principal line of argument against representationalism. This argument seeks to show that the representationalist cannot satisfactorily account for the relationship between perceptual representation and the fact that we see physical objects. There is, I argue, at least one sort of conception of this relationship that is available to the representationalist, independently motivated, and immune to Travis's objections.

⁶ This includes Travis's remarks on 'occasion-sensitivity' and its relevance to the proper philosophical view of perception. These remarks strike me as orthogonal to the issue of representationalism, since corresponding points seem to hold for beliefs and desires – the paradigm states with representational content. (See Richard (1990) for a discussion of context-sensitivity in belief reports.) But I will not be able to pursue this issue here.

Before we get to the main work of the paper, carried out in sections 3-6, the ground is cleared in the next two sections. In section 1 I discuss the four commitments Travis ascribes to representationalism. I go on to outline in section 2 Travis's account of how perceptual experience may be misleading, without involving misrepresentation. This part of his account is clearly designed to undercut the motivation for representationalism. It is notable also for the use it makes of his notion of a 'demonstrable look', a notion which turns out to be central to his case, and will occupy us at length in section 5.

1. The commitments of representationalism

Travis begins by setting out what he takes to be four crucial commitments of representationalism (he doesn't use that label for the view):

To sum up, the position on which in perception we are represented to, as I will construe it here, has four significant points.

1. The representation in question consists in representing things as so (thus, truly/veridically, or falsely/non-veridically).
2. It has, or gives, perceptual experience, a face value, at which it can be taken or declined (or discounted).
3. It is not autorepresentation. (It is allrepresentation, though here, not crucially.)
4. Where we are thus represented to, we can recognize that, and how, this is so; most pertinently, we can appreciate what it is that is thus represented to us as so. Provisionally, I suppose it is (in some sense) the way things look that lets us do that. (p. 63)

Point (1) here says that perceptual representation (were there such a thing) is (or would be) of a truth-evaluable kind.⁷ But it says more than that. Part of Travis's intention under point (1) is to distinguish perceptual representation from information carrying. (pp. 58-9) The configuration of rings in the trunk of a tree, say, may carry information that a certain summer was unusually hot. The information that a certain summer was unusually hot involves, on at least on some views, a truth-evaluable information content. Perceptual representation is different from mere information carrying, Travis suggests, in that it allows for the possibility of misrepresentation. Now, representationalists certainly ought to agree that perceptual representation is not mere information carrying. On the other hand, there is a certain strong way of understanding the claim that perceptual representation allows for the possibility of misrepresentation; on this strong understanding, representationalists are not obviously committed to it. On the strong understanding, to say that perceptual representation allows for misrepresentation is to say that a given perceptual experience has representational content only if that very experience could have that very content even if circumstances were such that the content were false. The claim that experiences have representational content does not commit one to this claim. At any rate, it does not so commit one if there are factive mental states such as knowing that P, seeing that P, and remembering that P.⁸ These factive mental states, if they exist, have representational content no less than beliefs or desires, but they do not allow for misrepresentation. A representationalist might hold there to be a core category of experiences that are factive mental states: that are of a kind that, by their very nature, can only have true representational content. Indeed, John McDowell and Bill Brewer have taken precisely that view.⁹ It would not

⁷ I will let the parenthetical qualifications here be tacit henceforth.

⁸ See Williamson (2000) for an extended argument for the existence of factive mental states.

⁹ See McDowell (1982: 386-387) and Brewer (1999: ch. 5)

be inconsistent for such a representationalist to claim that *only* experiences in this core category have representational content. On the other hand, it is fair to say that such a version of representationalism is a very much a minority view. For most representationalists, experiences (at least a central and widespread category of them) have both representational content and allow for misrepresentation, as understood on the strong requirement above.¹⁰ It would be extremely interesting if Travis succeeded in pointing out serious problems for any such view. At least *pro tem*, then, I will allow that representationalism endorses the possibility of misrepresentation, thus strongly understood.

Point (2) is fairly clear and ought, I think, to be an uncontroversial commitment of representationalism. The precise significance of point (3) is somewhat harder to pin down. Part of what Travis intends here is that perceptual representation is belief-independent: the fact that a perceptual experience has the content that P is not constituted by the subject's judging that P, or being disposed to judge that P, on its basis. Perhaps there is more to not being autorepresentational than being belief-independent. Nevertheless belief-independence seems capture at least the core of the notion. Not all representationalists have accepted belief-independence in this sense; indeed, early influential representationalists such as Armstrong (1968) and Pitcher (1970) explicitly took a belief-dependent view. However, belief-independence is independently plausible, and representationalists are certainly well advised to grant it.

Point (4) does most of the work in Travis's argument. We may call it the recognisability requirement. It raises many questions. Among other things, several representationalists

¹⁰ We shall return to the significance of the parenthetical qualification below.

will wonder who are included in the ‘we’, who need to be capable of recognising what the content of an arbitrary experience is. Is the idea that the subject herself needs to be capable of that, if we are to be entitled to say that her experience has representational content? This might seem too strong. Are we always capable of recognising what the content of our beliefs or desires are? Small children, as well as cats and dogs, plausibly have beliefs and desires; do we really want to say they can recognise the content of their beliefs and desires? Yet beliefs and desires are the paradigm states with representational content. Moreover, perhaps the majority of representationalists have held that perceptual experiences in some sense have *nonconceptual* representational content.¹¹ On the standard conception, this is taken to entail that the subject of an experience need not possess the concepts in terms of which we would specify the content of the experience. For example, in order to have an experience with the content that something has a certain highly specific shade of red, red_{345} say, a subject need not possess a concept of red_{345} . Recognition, on the other hand, may well be thought to require the deployment of concepts.¹² If so, and if nonconceptualism is right, we would precisely *not* expect subjects to be capable of recognising what representational content their experiences have. So one might think the ‘we’ would better be limited to third-party theorists of some sort who set their minds to it.

Travis clearly does not intend ‘we’ to be limited in any such way (p. 63). It is not clear to me how he would respond to the issues just raised. Perhaps he would say that recognising what content a given experience has would not take as much as one might

¹¹ Evans (1982), Peacocke (1992), Crane (1992), Tye (1995, 2000, 2006a, 2006b), McLaughlin (2003).

¹² Arguably, recognition of something as something is a case of propositional knowledge. Propositional knowledge plausibly entails belief, and belief, on the standard view, is conceptual. This indicates one sort of reason to think that recognition is conceptual.

presume; he does write that the recognition involved here ‘need not mean that we can characterize such representational content accurately, or formulate it explicitly’ (p. 62). In any event, and as far as I can see, it does not matter hugely to Travis’s arguments that the capacity for recognising what content their experiences have be one that all subjects could exercise forthwith. The critical claim for his purposes seems to be this.

(R) It must be possible to specify, in non-question-begging terms, facts about experiences that ground the non-arbitrary ascription of representational content, where these facts about experiences are available on reflection to normally self-conscious human adults, and, as per point (2), are non-autorepresentational.

The requirement that the specification be non-question-begging rules out using terms that obviously, or by stipulation, is bound up with representational content talk. We cannot, for example, specify the sought-after facts in terms of what the experience is ‘intentionally directed upon as being the case’. It also rules out specifications in terms of ‘veridicality’ or ‘correctness’ conditions, since the applicability of these terms is very much part of what is in question. The requirement that the sought-after facts be available on reflection to normally self-conscious subjects ensures that they are facts about *conscious* experiences as such. After all, representationalists don’t claim that only *subconscious* features of perceptual states can underwrite the ascription of representational content. Interpreted as saying (R), representationalists ought, I think, to be able to accept the recognisability requirement.

2 Misleading but not misrepresenting

The idea that perceptual experiences admit of a principled classification as veridical on the one hand, or illusory or hallucinatory – for short delusive – on the other is absolutely

central to representationalism. Following Austin (1962: 11), and indeed Descartes (1971: 96-7), Travis holds that the notions of veridicality and delusiveness at best have a metaphorical application to sense experience. He admits, as one surely must, that the senses may mislead, but argues that misleading sensory experiences in no way need to be explained in terms of misrepresentation.

His model for accounting for misleading sensory experiences is misleading facts. Consider the fact that Jones's car is in the driveway. Typically, this means that Jones is at home. The meaning in question here is what Travis calls 'factive meaning'. That A factively means that B only if one can soundly reason: 'Since A, B'. For example, if the fact that Jones's car is in the driveway factively means that he is at home, one can soundly reason: 'Since Jones's car is in the driveway, he is at home.' Notice that one can soundly reason in this way only if Jones's car is in the driveway and he is indeed at home. So factive meaning is by nature never misleading. However, the fact that A can mislead if it 'might reasonably, or rightly, be expected factively to mean B' (p. 67), that is, in Travis's term, if it 'indicates' that B, even though it is not the case that B. Thus, if, on this exceptional occasion, Jones's car is in the driveway but he is not at home, the fact that his car is in the driveway is misleading, in that it indicates something that is not the case.

Travis exploits this model to give the following simple account of misleading visual experiences. Visual experiences mislead when what one sees indicates something that is not the case. Thus if I see the rear half of a pig sticking out behind a corner, what I see indicates that there is a pig in that spot. Yet if there is *only* a rear half of a pig in that spot, cleverly set up to make it look as though the rest of the pig was half-obscured by

the wall, the presence of the rear half does not factively mean what one might reasonably expect it to mean. The misleading character of my visual experience is due to my seeing something which, in turn, indicates something that is not the case.

The application of this model is most straightforward when it is fairly clear what the seen ‘mis-indicating’ object or state of affairs is. In the above case, it is pretty clear that the presence of the rear half of the pig was *both* something that mis-indicated the presence of a pig, *and* something seen. The application of the model is perhaps less straightforward when it comes to what seem like simpler or more basic cases of a visual illusion, such as the Müller-Lyer. What is it that makes this experience misleadingly as of two lines of unequal length? An answer that does *not* work here is to say that this is the arrangement of the equally long lines on the page having the characteristic ‘arrows’ attached to them. This arrangement is certainly something we see, but is not something that indicates that the lines are unequally long. The fact that two lines on a page have such-and-such wedges attached to them might not reasonably, or rightly, be expected factively to mean that the lines are unequally long. Yet if we cannot cite the arrangement of the lines and wedges on the page as an answer, one might wonder what one can cite as an answer. There are no obvious, more proximate objects we see than the arrangement of the lines on the page, or so one might think.

Travis, somewhat strikingly, handles this apparent difficulty in terms of his notion of a ‘demonstrable look’. This notion of a demonstrable look is expressed by using ‘look’ as a count noun, taking determiners such as ‘a’ and ‘the’. Sometimes (but not always) Travis highlights what he takes to be a significant feature of the notion, namely that the looks in question are things we can see and think of demonstratively, as *that* look, by

adding ‘demonstrable’ in front the count noun use of ‘look’. On the Müller-Lyer, he writes,

In the Müller-Lyer, two lines are contrived (by means of accompanying wedges) to have a certain look. They do not just *seem* to have that look; that is actually the way they look. Two lines may well have that look because one *is* longer than the other. That is a familiar way for things to be. Depending on circumstances, that look may thus *indicate* that it is two lines of unequal length that one confronts. ... False expectations arise here in the wrong view of what something (a look) means, though perhaps a right view of what it *ought* to. What one gets wrong is the arrangement of the world: how the misleading seen thing in fact relates to other. (p. 68)

One thing that is somewhat striking about this account is the strict parallels it draws between the ‘detached-half-of-a-pig’ illusion and the Müller-Lyer case. Another somewhat striking feature is the way it talks of demonstrable looks as things we see: as objects of vision.¹³ How are these demonstrable looks to be understood? Isn’t the role for which they are here cast rather like that played by sense-data on sense-datum theories? If so, wouldn’t that be worrying – so worrying, perhaps, that an account of illusion in terms of misrepresentation would be preferable after all? We shall come back to these issues below.

A third striking feature by the account is brought out by the following question: why does Travis give the account in terms of demonstrable looks rather than what might seem to be a simpler account in the vicinity? The apparently simpler account says that the two lines in the Müller-Lyer drawing look unequally long to the subject. The fact that two lines look unequally long to one might (depending on the circumstances)

¹³ Travis makes explicit this role of demonstrable looks at various places: ‘A demonstrable look, like any other visible thing, may sometimes indicate something. But that is a contingent matter’ (p. 80).

reasonably be expected to mean that they are unequally long: it indicates that they are unequally long, although in this case misleadingly. This is what makes the visual experience misleading, for the visual experience consists, in part, in the two lines looking unequally long to one. The present account makes no reference to demonstrable looks as things that we see in the M-L illusion. True, it has the consequence that the structure of the M-L case is not quite parallel to that of the ‘detached-half-of-a-pig’ illusion, but it is not clear why that should count against the account rather than in its favour, or be considered irrelevant.

Perhaps Travis does not think of his account as in any way significantly different from the apparently simpler one just given. Or perhaps he thinks his account gives a more lucid view of matters described by our apparently simpler account. I believe the latter is a reasonable interpretative conjecture. We shall come back to these issues in section 5 below, where we will examine the notion of demonstrable looks more closely.

3. Looks-indexing I: looking a certain way

Travis provisionally supposes that it is (in some sense) the way things look, in an experience, that lets us recognise what is represented to us as so, in that experience. This amounts to supposing, in his terms, that representational content is ‘looks-indexed’. In light of our interpretation of the recognisability requirement in terms of (R), this comes to the following supposition

Looks-indexing A specification of the ways thing look, or appear, constitutes a specification, in non-question-begging terms, of facts about experiences that ground the non-arbitrary ascription of representational content, where these facts are

available on reflection to normally self-conscious human adults, and are non-autorepresentational.

Travis does not maintain that the representationalist is committed to looks-indexing. He distinguishes three alternative ways in which the representationalist might seek to meet the requirements of (R). The first alternative is that the ascription of representational content is grounded on facts about what we see; the second that it is grounded on facts about what we see* (pp. 85-7); the third, on fact about what we ostensibly see (pp. 79-81).¹⁴ I agree with Travis, however, that these options are no more promising than looks-indexing, and fail if look-indexing does. Travis does not argue that the four options thus distinguish exhaust the representationalists' options for meeting the demands of meeting (R); he could fairly point out, however, that if none of these four options work, the onus is certainly on the representationalist to come up with something that does. I believe the representationalist ought to be prepared to take on board a commitment to looks-indexing.

The burden of Travis's paper is devoted to showing that representational content is not looks-indexed. There are two notions of looks relevant to looks-indexing, he argues, and, on neither of these, facts about how things look can ground the ascription of representational content in the sought-after way. Travis's second notion is that characteristically expressed by such remarks as 'It looks as if Pia will sink that putt', or 'It looks like Frank will win the race.' Travis takes the notion of looks thus expressed to be one on which its looking as if P to S is to be elucidated in terms of there being

¹⁴ 'See*' is to express an intensional reading of 'see', which Travis finds in Harman (1990). On this reading, we can truly say 'Sally sees a tree' even if we cannot truly say 'There is a tree and Sally sees it'. The term 'ostensibly sees' is one Travis takes from McDowell (1998a). In the end, he finds that 'seeing*' and 'ostensibly seeing' effectively amount to the same notion (p. 87).

certain body of facts (typically including facts about how things look, on his *first* notion of looks, to which we will turn in a moment) that indicate or factively mean that P, and, on the basis of which S takes it that P, perhaps tentatively (pp. 77-8). Travis's second notion of looks is thus roughly equivalent to Jackson's (1977: 30-1) conception of what he calls the 'epistemic use' of 'looks'.¹⁵ Facts about how things look, on this second, epistemic notion, cannot ground representational content, Travis argues, because neither indication nor factive meaning makes for representation, and because they are sensitive to how people take things to be in ways in which properly perceptual representation is not. Whether for this reason or for others, I believe we ought to agree that the second, epistemic notion of looks is a bad candidate for vindicating looks-indexing. For the rest of the paper, I set it aside.

Travis introduces his first notion of looks as follows:

On the first notion, something looks thus-and-so, or like such-and-such, where it looks the way such-and-such, or things which are (were) thus and so, does (would, might) look. On this notion, Pia may look (rather, very much, exactly) like (the spitting image of) her sister. (She and her sister look alike; she looks the (same) as her sister does.) That man on the bench looks old. (He looks the way an old man would, or might.) The shirt looks blue (in this light) – as a blue shirt (so viewed) does, or might. (pp. 69-70)

This notion of looks cannot determine any representational content for experiences, Travis argues, because, if it points in any direction at all, it points in too many, mutually incompatible, directions. That is to say: if a specification of how things look to one, in terms of this first notion, grounds the ascription of any representational content, it

¹⁵ A not very important contrast is that Jackson, unlike Travis, allows that it is strictly coherent to say 'It looks as if P, but I happen to know that not P'.

grounds the ascription of an incoherent, unsatisfiable content. But it is absurd to think that any perceptual experience has an incoherent content. So the relevant specification does not ground the ascription of representational content after all.

The critical premise of this ‘incoherence argument’ is what he may call the ‘incoherence claim’.

Incoherence claim If a specification of how things look to one, in a given visual experience, in terms of this first notion, grounds the ascription of any representational content to it, it grounds the ascription of an incoherent content to it.

Travis gives the following recipe for how we might illustrate the incoherence claim, for an arbitrary experience:

Take any way things may be said to look. Now take any way that things may fail to be what they would need to be to be what they *thus* look like. That is *another* way things may be said then to look: they look just the way they would if *that*, rather than the first things, were the way things are. ... [T]he obtaining of this second thing, or of anything else that made things fail to be that first thing, would make for misrepresentation only if something decided that it was only the first fact about looks, and not the second, or any other such, that indexed representational content. Perception does no such selecting for us. (p. 74)

We may flesh this out in terms of an example Travis himself uses (pp. 73-4). Say that you are looking at a shirt, and the shirt looks blue to you. The way the shirt would have to be to be what it *thus* looks like is blue. (That, at any rate, is Travis’s idea if I understand his turn of phrase correctly.) So, if the fact that the shirt looks blue to you

determines any representational content (or part thereof) for the experience, it determines the content that the shirt is blue. Now consider a scenario (which may or may not be actual) on which the shirt is not blue yet, as we say, looks just as it actually does. The case in which the shirt is white but cleverly illuminated by blue light is such a scenario, it seems. Since the shirt looks the same in this scenario, our shirt, in looking blue, looks the way white things cleverly illuminated by blue light do (would, might) look. That is to say that, in looking blue, our shirt looks like a white shirt cleverly illuminated by blue light. The way the shirt would have to be to be what it *thus* looks like, Travis suggests, is a white shirt illuminated in blue light. In other words, if the fact that the shirt looks like a white thing in blue light determines any representational content (or part thereof) for the experience, it determines the content that the shirt is a white shirt illuminated by blue light. It follows then, that if any content gets determined for the experience, the content that the shirt is blue, and white and illuminated by blue light, gets determined.¹⁶ But nothing can be both blue and white. The content is unsatisfiable, incoherent. Travis anticipates the response that only one of these looks-facts, the fact that the shirt looks blue, is genuinely content-determining. He maintains, however, that there is that there is nothing about perception that underlies this selection.

Let's set aside for a moment the question of whether Travis is right that both looks-facts are content determining if either is. Another feature of his argument that ought to attract our attention is that it sets no store by a small but possibly significant difference between the two looks-facts in play. We started out with the fact that the shirt looks blue. It was argued that this entailed the fact that it looks *like* a white thing illuminated

¹⁶ Alternatively, counting things differently, both the content that the shirt is blue, and the content that the shirt is white and illuminated by blue light, get determined as contents of the experience, perhaps as parts of its overall content. It does not matter if we look at it this way or as in the text.

by blue light. The former looks-fact was supposed to determine the content that the shirt *is* blue, the latter, likewise, the content that it *is* white and illuminated by blue light. The presence of the word ‘like’ is not supposed to matter, then, to the content-determining force of the looks-fact. This is not irrelevant to the argument, even if we agree that both facts are content determining. What if the content determined by our second fact was, rather, that the shirt is *like* a white thing illuminated by blue light? Then we get no incoherence when we conjoin the contents. Something may well be blue and *like* a white thing illuminated by blue light.

This disappearance of the word ‘like’, in the transition from the specification of the looks-fact to the specification of the putative content or veridicality condition, is systematic in Travis’s illustrations of the putative incoherence:

If Pia looks like her sister, there is a way she should be to be what she *thus* looks like: she should be her sister. If Pia’s imitation of her sister nonplussed looks uncannily like her sister nonplussed, then, again there is a way things should be to be what they *thus* look like: it should be her sister, nonplussed. (p. 70)

In looking like her sister Pia shares a look with countless other things – herself, a wax replica of herself in Madame Tussaud’s, a good hologram. ... For each of these, there is a way things should be to be what they *thus* look like Pia should be, respectively, herself, a wax dummy, a hologram, ... (p. 72)

An obvious representationalist rejoinder to Travis is that these disappearances reflect a fallacy. The rejoinder might be developed as follows. Let’s call looks-facts, or statements of such facts, *qualitative* iff they are articulated on the form ‘*x* looks F’,

where ‘F’ is an adjective, for example ‘red’, ‘round’, ‘happy’, or ‘interesting’.¹⁷ Let’s call looks-facts, or statements of such facts, *explicitly comparative* iff they are articulated on the form ‘x looks like so-and-so’.¹⁸ Some qualitative looks-statements may be *implicitly* comparative, in that they properly may be understood in terms of, and have the same truth-conditions as, a corresponding explicitly comparative statement. For example, one might think that the statement that x looks American properly may be understood in terms of, and has the same truth-conditions as the statement that x looks like an American. I will refer to qualitative looks-statements that are *not* implicitly comparative in this way as *non-comparative* qualitative. The central claim of the representationalist rejoinder is this:

Non-comparative Grounding For any visual experience, (i) there is a class of non-comparative qualitative looks-statements true of it, and (ii) the non-comparative qualitative looks-facts about the experience affirmed by these statements ground the ascription of at least the core of its representational content.

The representationalist may take this key idea further in different ways. She may challenge Travis head-on and claim that only non-comparative qualitative looks-facts are content determining. Alternatively, she may allow that explicitly or implicitly

¹⁷ For present purposes, facts can be equated with true propositions. I will alternate between talking in terms of looks-facts and looks-statements.

¹⁸ Travis notes that there is a use of ‘looking like’ to express his second, epistemic notion of looks. He plausibly notes, moreover, that ‘looking like’ is *constrained* to express this second, epistemic notion when it is complemented by a sentence in the indicative. For example, ‘It looks like that person is an American’ is constrained to be understood on this second, epistemic notion, as saying roughly that, going by the way things look (in a non-epistemic sense of ‘looks’) it probably means that that person is an American. An utterance that ‘x looks like an American’ may be interpretable as expressing this epistemic notion too, or so Travis suggests. Comparative looks-statements are however, by definition, not to be so interpreted.

comparative looks-facts are also content determining, but insist that these facts, unlike non-comparative qualitative facts, determine only comparative contents. For example, that Pia looks like her sister, like a wax dummy, and like a hologram determines the content that Pia *is like* her sister, like a wax dummy, and like hologram. Perhaps it will be protested that the latter content surely will not be true, even if it is not strictly incoherent, and that the content, for that reason, is not much better than an incoherent content. However, the representationalist maintains that there is at least one good sense in which Pia *is like* each of these things; among other things, she is visually like each these things. Indeed, for each of them, there may be countless ways in which Pia is like it. Admittedly, these comparative contents are not very informative. But they do not make for incoherence, or raise the standard of veridicality for the experience unreasonably high.

Travis, it seems, would not accept the idea of there being a class of non-comparative qualitative looks-statements. When he introduces his first notion of looks, he writes in wholly general terms:

On the first notion, something looks thus-and-so, or like such-and-such, where it looks the way such-and-such, or things which are (were) thus and so, does (would, might) look. (pp. 69-70)

He here seems to explain in the same breath what it is for something to look blue, say, and what it for something to look like a blue thing. If he is right in doing so for any qualitative looks-statement, one would expect any fact of the form '*x* looks F', and the corresponding explicitly comparative fact of the form '*x* looks like an F thing', to be on a par when it comes to content determination. That result would certainly refute the

head-on challenge to Travis that a core class of qualitative looks-facts is content determining while no comparative facts are. It would also refute Non-comparative Grounding as stated above.

The result would invalidate moreover the more general representationalist idea that there is at least some contrast between qualitative and comparative looks-facts when it comes to content determination, even if both determine content. Suppose, as this more concessive response allows, that x 's looking like a blue thing only determines the content that x is like a blue thing. The fact of x 's looking blue would, then, only determine that comparative content too. Since there is an explicitly comparative looks-fact corresponding to each qualitative looks-fact, no looks-facts determine anything more than comparative contents.¹⁹ Admittedly, this does not make the content incoherent. But it makes it well-nigh trivial, since anything is like anything (in some respect); surely an untenable result. If, on the other hand, we want to allow x 's looking blue to determine the content that x is blue, we must allow that the corresponding explicitly comparative fact determines that content too. Yet then we would similarly be committed to allow that the shirt's looking like a white thing in blue light determines the content that it is a white thing in blue light. The incoherence claim is thereby reinstated.

It is apparent, then, that Travis's incoherence argument brings to the surface an underlying question of principle: the question, namely, of what the semantic or metaphysical interconnections between qualitative and explicitly comparative looks-

¹⁹ I am assuming here, for simplicity, that qualitative and explicitly comparative looks-facts exhaust the looks-facts. This assumption is not strictly correct, of course, but we can safely set other notions of looks aside for the purposes of discussing the incoherence argument.

statements may be. Chisholm (1957) and Jackson (1977) both devoted considerable attention to this question, arguing, in effect, that there is a core class of non-comparative qualitative looks-statements. In the next section, I shall argue that they were right about this.

Apart from this, our response, on Travis's behalf, to the representationalist rejoinder suffers from another drawback. The response rightly points out that for each qualitative looks-fact there is a corresponding explicitly comparative fact. Yet the converse is not obviously true. Consider such looks-facts as

- (1) The shirt looks like a white thing illuminated by blue light.
- (2) Pia looks like her sister.
- (3) The creature looks like a centaur.

What would be the qualitative counterparts of these? Certainly not

- (1a)* The shirt looks white thing illuminated by blue light
- (2a)* Pia looks her sister.
- (3a)* The creature looks a centaur.

These are simply ill-formed. Now, if (1)-(3) do not have qualitative counterparts, one might think this is significant for the purposes of the incoherence argument. For one might think that *even if* one has to allow that *x*'s looking like a blue thing determines the content that *x is* blue (not merely the that *x is like* a blue thing) because of the presence of the corresponding qualitative looks-fact, we do not have *that* reason to allow that the shirt's looking like a white, bluishly illuminated thing, or Pia's looking like her sister, determines the content that the shirt *is* a white bluishly illuminated thing, or that Pia *is* her sister. This drawback to our response, on Travis's behalf, may not be fatal. But it

shows that even given the parity between x 's looking blue and x 's looking like a blue thing, there is work to be done in re-instating the incoherence claim.²⁰

Travis, we took it, would object to our representationalist rejoinder on grounds of its commitment to a class of non-comparative qualitative looks-statements. However, there is another interesting objection to our representationalist, that concedes the existence of this class, and even that only statements within it determine representational content. The objection is that, even conceding all that, the incoherence argument can be revived.²¹ Consider again the shirt looking blue to us. Travis argues (and we did not object) that it *ipso facto* looks like a white thing illuminated by blue light. But another, closely related way it looks is this,

(4) The shirt looks white and illuminated by blue light.

The fact that the shirt looks as specified by (4) determines the content that the shirt is white and illuminated by blue light. At any rate, this is certainly what we should expect if the fact that the shirt looks blue determines the content that it is blue. It follows that

²⁰ One way a defender of Travis might try to deal with this difficulty is as follows. There is a way of finding a qualitative statement corresponding to each explicitly comparative statement, for we have, or can legitimately devise, a generally applicable way of finding an adjective corresponding noun phrase. English has various devices that are applicable to some nouns for this purpose, viz. applying the -ean suffix. Thus corresponding to 'antipode' we have 'antipodean'. It might not be too much of a stretch to conceive of this device as being more generally applicable than it in fact is. Thus, corresponding to 'centaur' we have 'centaurean', to 'her sister' we have 'auto-sororean', to 'white thing illuminated by blue light' we have 'white-illuminated-by-blue-light-ean', and so on. This gives us as counterparts to (1)-(3):

- (1b) The shirt looks white-illuminated-by-blue-light-ean
- (2b) Pia looks auto-sororean.
- (3b) The creature looks centaurean.

I am not sure if Travis would want to endorse this line of thought. Moreover, it invites the objection that our understanding of what it is to look centaurean is wholly provided by our understanding of what it is to look like a centaur, whereas, by contrast, there is not the same reason to think that our understanding of what it is to look blue is wholly provided by our understanding of what it is to look like a blue thing. Jackson (1977: 88-89) makes effectively this objection to Chisholm's (1957: 115) deployment of sentences like (3b) against the sense-datum theory.

²¹ Herman Cappelen pressed this objection (in conversation).

we have once again an incoherent overall content on our hands: nothing can be both blue and white (and illuminated by blue light).

This objection fails, I maintain, since the shirt, in looking blue, does not look white and illuminated by blue light. This is because ‘looks’ distributes over conjunction (at least) in qualitative looks-statements. Thus (4) entails

(5) The shirt looks white and looks illuminated by blue light.

Which of course entails

(6) The shirt looks white

But this is clearly false. The shirt does not, in looking blue, look white.²²

To sum up before we move on. Travis’s critical argument against looks-indexing, in terms of his first notion of looks, is the incoherence argument. The critical issue underlying this argument is, in turn, whether qualitative looks-statements and corresponding explicitly comparative looks statements are on a par when it comes to content determination. Travis seems to suppose that they are; indeed, this seems to be built into his first notion of looks, the ‘looking-like’ notion. So one way of putting the issue at stake here is whether Travis fails to recognise a notion of looks: the non-comparative qualitative one. If there is non-comparative qualitative notion Travis fails to recognise here, it is precisely the one that is the representationalist’s best bet to validate looks-indexing.

²² Of course, it might look blue to some and white to others. We could have avoided this possibility by specifying a constant subject throughout. Other contextual parameters apart from subject are also implicitly to be held constant.

4. Are some qualitative looks-statements to be understood non-comparatively?

In Chisholm's (1957) and Jackson's (1977) classic works on perception, considerable attention is devoted to distinguishing and comparing various constructions involving 'looks' and 'appears'. Chisholm and Jackson allow that some qualitative looks-statements involve a use of 'looks' that is comparative. Chisholm writes, for example,

Using appear words in the present nonepistemic sense, we may say that the railroad tracks 'look convergent' if they look the way two converging lines would ordinarily be expected to look if both were visible at once. ... [T]he point of the locution 'x appears so-and-so', in its present use, is to compare *x* with things that *are* so-and-so. Let us speak, then, of the *comparative* use of appear words. (Chisholm, 1957: 45)

On Chisholm's account, then, the relevant fact that the railway lines look convergent, as stated by the relevant use of 'look convergent', is implicitly comparative. Chisholm's explanation of what 'x looks convergent' states, on this use, is clearly related to the account we get if we apply Travis's schema (pp. 69-70), viz. that two things look convergent when they look the way convergent things do (would, might) look. The differences between Chisholm's account and Travis's do not seem to matter to our purposes; we may classify them together as accounts on which the looks-statement is understood comparatively.²³

Chisholm argues, however, that not all qualitative looks-statements are to be understood comparatively:

²³ I will, however, spell out the comparative understanding according to Travis's schema on pp. 69-70 in what follows, since he is our target here.

When 'looks red' is used comparatively, the statement

(3) Things which are red look red in daylight.

is analytic, for it says, of the way red things look in daylight, only that it is the way red things look in daylight. ... But when 'looks red' is taken noncomparatively, (3) is synthetic – an 'empirical generalisation.' (1957: 51)

One way of developing Chisholm's line of thought here is as follows.²⁴ It is a mistake always to take

(7) F things generally look F.

to mean the same as

(7a) F things generally look the way F things (would, might) look.²⁵

This is a mistake since there are uses of instances of (7) that are false for certain circumstances for which the corresponding instance of (7a) is either true or false for the wrong reasons (depending on how we take the Travisian parenthesis). Consider the pair

(8) Isosceles things generally look isosceles.

(8a) Isosceles things generally look the way isosceles things (would, might) look.

Suppose that, as a result of some cosmic ray affecting human shape perception in hard-to-reverse ways, isosceles things generally look somewhat scalene. There is a perfectly normal use of (8) on which it is false, relative to this scenario. What are we to say about (8a) relative to it? In the scenario, isosceles things certainly generally look the way isosceles things look. Since the imagined change in shape perception is hard to reverse, isosceles things would, moreover, look as they do in the target scenario in all possible circumstances that are close by our target. So it seems to be true, relative to our

²⁴ See Jackson (1977: 33-37) for closely related arguments. What Jackson calls 'phenomenal uses of looks' are a proper subclass of what we have been calling qualitative looks-statements.

²⁵ The parenthesis here reflects the parenthesis in Travis's account of his first notion of looks on pp. 69-70, quoted above.

scenario, that isosceles things look the way isosceles things do or would look. Perhaps someone would argue that isosceles things do not, in the envisaged scenario, generally look the way they *might* look. For they might look isosceles (after all, they actually do). However, if this is the reason why (8a) is false, it seems to be false for the wrong reason. Isosceles might also look scalene. So by parity of reasoning, we get the incoherent result that (8a) is also true under the envisaged scenario. Under the envisaged scenario, then (8a) is either true, or incoherent.

A defender of Travis might be inclined to deal with this problem by rigidifying (7a), so that it says that F things generally look the way F things *actually* look, or, perhaps, look in circumstances close to the *actual* world. Yet difficulties for the comparative construal of qualitative statements arise also for actual cases. Mark Johnston (2004: 141-2) has recently drawn attention to the intriguing shade of colour known as ‘supersaturated red’. This is a shade of colour which one might, as a matter of a law about our visual system, experience only in the context of an afterimage-involving experience. To get an experience of supersaturated red one needs to be exposed for about twenty minutes to bright monochromatic unique green light in an otherwise dark room. When then lights are turned on, one will experience a red afterimage. If one orients one’s gaze in such a manner that the red afterimage is superimposed on a small red background, one will be having an experience of supersaturated red, a red more saturated than any surface or spectral red we can see. Now take such a case where we have that

(9) Things look supersaturated red to S.

To simplify, let’s assume that (9) describes S’s entire visual field at a given time. The question to ask now is whether (9) can be spelled out in terms of any of the following:

(9a) Things look to S the way supersaturated red things do.

(9b) Things look to S the way supersaturated red things would.

(9c) Things look to S the way supersaturated red things might.

It would not be unreasonable to think that nothing is supersaturated red.²⁶ If there are no supersaturated red things, there is no particular way they look. It follows from (9a), then, that things look no particular way to S. But this is false; at the very least things look some shade of red to S.

Even if there are supersaturated red things somewhere in the universe, (9a) is either false or true for the wrong reasons. Note first that that these supersaturated red things (wherever they may be) do not tend to look, nor typically look, nor, in standard conditions, look supersaturated red. After all, situations in which the subject has an afterimage in the line of sight of the object of a given kind are not relevant to determining how things of that kind tend to look, or typically look, or look in standard conditions. The laws of our visual system being as they are, it is only in those situations that things look supersaturated red to us. Perhaps it will be objected that, even if afterimage experiences are irrelevant in this way as a general rule to how things tend to look, we need to make a special exception when it comes to the question of how supersaturated red things tend to look. This objection, however, seems entirely *ad hoc*. So if the way supersaturated red things look, for the purposes of (9a), is the way they tend to look, or typically look, or look in standard conditions, (9a) is false.

²⁶ Perhaps it will be objected: but is not the afterimage itself supersaturated red? To this one might respond, first, that it can plausibly be argued that, strictly and literally, there are no afterimages. What there are are afterimage-type visual experiences. Talk of ‘having afterimages’ or ‘seeing an afterimage’, to the extent that it says something true, is to be interpreted as a *façon de parler*. Second, even if there were afterimages, they would not be supersaturated red for the same reason as they would not be red, or blue. The predicates ‘...is red’, or ‘... is blue’, apply to physical surfaces, or volumes, or light sources, if they apply to anything at all. See Peacocke (1983: ch. 2; 1984) for closely related arguments.

Perhaps it will be objected, at this point, that, even if supersaturated red things do not tend to look, nor typically look, nor in standard conditions look supersaturated red, (9a) is still true. If this is so, (9a) is true only because supersaturated things in exceptional conditions look supersaturated red. But this condition is not sufficient for the truth of (9). For all we have said, supersaturated red things in exceptional conditions look blue. But it does not follow that things look blue to S. Either way, (9) and (9a) diverge in truth-conditions.

Similar remarks extend to (9b). 'x looks the way F things would look' is reasonably understood to say that x looks the way F things would look, the laws of our visual system being as they are. For example, when we wonder how sailboats made of gold would look out on the water we are wondering how sailboats made of gold would look out on the water, the laws of our visual system being as they are. If they were different, who's to say what they would look like! We can, then, run the argument against (9a) *mutatis mutandis* against (9b). Consider a scenario in which the laws of our visual system are as they are. Any supersaturated red things there are in that scenario will not tend to look, not typically look, nor, in standard conditions, look supersaturated red to a subject like S. So, taken relative to that scenario, (9a) will be either false or true for the wrong reasons. It follows that the ways supersaturated red things would look does not make (9b) true.

Nor does (9c) plausibly spell out (9). Let's grant that in some remote possible scenario there are supersaturated red things and they tend to look, or typically look, or in standard conditions look supersaturated red to subjects like S. Perhaps granting so much is sufficient to make (9c) true, at least given that something actually looks some way to

S. However, granting this much is not sufficient to make (9) true, as the following analogy shows. There is a remote scenario in which happy creatures tend to look pyramid-shaped. The existence of this scenario is sufficient to make it true, in a sense, that the Cheops pyramid looks a way happy creatures *might*. Yet it would be silly to conclude that the Cheops therefore looks happy.

These apparent problems of understanding (9) in such comparative terms as (9a)-(9c) points to an underlying problem about the idea that all qualitative looks statements are susceptible of a comparative understanding. The idea has the consequence that the *way* something looks always, on the more fundamental, explanatorily prior level, is to be picked out indirectly, via restrictions on the kinds of objects that look some way to us, and on the circumstances in which they do so. The relevant *way* things look is then specified as being the way objects that are *such-and-such* look when *so-and-so*. Reflecting on the case of supersaturated red suggests that this order of conception gets things the wrong way around. Surely, the more fundamental conception of the way things look when they look supersaturated red to us is the one visually available to us when enjoying an experience in which something looks supersaturated red to us. I may then concentrate on the way things look to me, colour-wise, and say to myself, 'That shade of colour, which these things now look, and which, for all I know, these things may in fact have, that shade is supersaturated red.' The conception I thereby get of that shade is one which I may, in turn, use to frame such questions as, 'I wonder what way a wallpaper with that shade of colour would look in the living room?' Suppose that, contrary to fact, the way a wallpaper with that shade would look in the living room would be supersaturated red. In that case, our conception of the way a wallpaper with that colour would look in the living room is a conception of supersaturated red. Yet

clearly that conception is less fundamental than my visual conception of supersaturated red.²⁷

I conclude that at least some qualitative looks-statements are to be understood noncomparatively. This class plausibly includes some statements to the effect that *x* looks blue. The corresponding explicitly comparative looks-statement, *x* looks like a blue thing, may of course be understood comparatively. Thus some corresponding qualitative and explicitly comparative looks-statements are semantically distinct. Our representationalist rejoinder to the incoherence argument was right in assuming the existence of a class of non-comparative qualitative looks-statements.

5. Looks-indexing II: Demonstrable looks

The notion of a demonstrable look looms large in Travis's overall argument. Among other things, it has a central role in his account of how visual experience can be misleading, as noted above. The notion is evidently supposed to contrast with his second, epistemic notion of looks, the notion characteristically expressed by such constructions as 'It looks as if Lisa will sink the putt' and 'It looks like Peter will win the race.' Since Travis early on announces that he will be 'distinguishing, and exploring, two different notions of looks' (p. 69), this might lead one to expect that the notion of a demonstrable look is meant to be somehow equivalent with, or a subspecies of, what he refers to as his first notion of looks, i.e. the notion according to which 'something looks thus-and-so, or like such-and-such, where it looks the way such-and-such, or things which are (were) thus and so, does (would, might) look.' To be sure, his

²⁷ For an interesting discussion of how visual experience, even in illusion or hallucination, is an original source of the possibility of *de re* thought about qualities and relations, see Johnston (2004: 129-134).

discussion of demonstrable looks is very closely bound up with his discussion of his first notion of looks. At the same time, the way he phrases things sometimes suggests, ever so mildly, a contrast between the notions. He writes, for example, that ‘no demonstrable look, *nor* any look in our first sense, *has* any intrinsic import.’ (p. 81, first emphasis mine)

Moreover, there is at least one significant claim of his that would seem to distinguish between the notions. Travis takes it that demonstrable looks are things that are seen. For example, when I look at the Müller-Lyer set-up, I *see* a certain demonstrable look, which is the look of the arrangement of the lines on the page.²⁸ If I am in a situation where a shirt looks blue to me, or Pia looks like her sister to me, it does not seem natural to say that I *see* the shirt’s looking blue to me, or Pia’s looking like her sister to me. True, there may be a more or less natural sense in which I can say that I see *that* the shirt looks blue to me, or that Pia looks like her sister to me.²⁹ Yet the sense of ‘see

²⁸ Cf. fn. 13 above. To say that I see this demonstrable look is clearly not, in Travis’s mind, intended to imply that I do not see the lines on the page. Nor is it intended to imply, I take it, that my seeing of the lines on the page is somehow indirect. We shall return to whether these intentions can be fulfilled.

²⁹ At this point, it might be useful to note a possible source of confusion over such sentences as ‘The shirt looks blue to Tim’. There is a reading of such sentences in which it does not imply that Tim sees the shirt (at the moment at which the sentence is uttered). This reading of the sentence is analogous to that on which ‘Sarah drinks Coca-Cola’ can be true even though Sarah does not drink anything at the time of utterance, and ‘Peter receives a lot of attention’ can be true even though everyone who has heard of Peter happens to be in a dreamless sleep at the moment of utterance. These readings are sometimes referred to as generic, and thought to have the character of attributing dispositions. On the generic reading, ‘Sarah drinks Coca-Cola’ says, in effect, that Sarah is disposed to drink Coca-Cola in such-and-such tacitly understood circumstances. The reading on which ‘The shirt looks blue to Tim’ does not entail that Tim sees the shirt at the moment of utterance seems to be generic in this way: it seems to say, in effect, that the shirt is disposed to look blue to Tim in such and such tacitly understood circumstances. On such a reading of ‘looks’, the shirt’s looking blue to Tim is a non-starter for constituting, even in part, a visual experience enjoyed by Tim, since it is consistent with Tim sleeping dreamlessly. To appreciate the suggestion that the shirt’s looking blue to Tim might constitute part of the visual experience Tim enjoys – a suggestion I take to be plausible – it is important, then, to set this generic reading aside. For more on generics, see Cohen (1999).

that' that makes this a more-or-less natural thing to say is precisely the sense in play in such clearly non-visual contexts as 'I see that the PM is about to resign' and 'I see that the conclusion follows'. It may also be that we can conceive of eccentric situations in which one could truly say that I see the event of the shirt's looking blue to me.³⁰ But the visual situations in which Travis intends to be so that I see a demonstrable look are generally not eccentric in this way. So even if I could in some way see the shirt's looking blue to me, this seeing would appear to be a different kettle of fish than the seeing that is intended to be the seeing of the demonstrable look of the shirt.

Let's call the fact that such-and-such a thing has a certain demonstrable look a 'demonstrable look fact'.³¹ The above remarks suggest, at the very least, that demonstrable look-facts are non-trivially related to both qualitative and explicitly comparative looks-facts. Again, it is not entirely clear to me what Travis takes the non-trivial relation to be, nor, indeed, whether he thinks of it as non-trivial. In any case, the following is clear. If the relation were trivial, the notion of demonstrable looks wouldn't add anything to the argument couched in terms of qualitative and explicitly comparative looks-facts. Since the latter argument – the incoherence argument, as stated above – has been found wanting, that would be bad for Travis's case. Further, if demonstrable looks talk is supposed to be a *façon de parler* to be cashed out in terms of qualitative or comparative looks talk, or if demonstrable looks-facts are supposed to be susceptible of a deeper, more illuminating characterisation in terms of either qualitative or explicitly

³⁰ For example, a physicalist might think that the event of the shirt's looking blue to me is identical to a complex physical event involving light bouncing off the shirt, reaching my eyes, and pulling into play various parts of my visual system. Perhaps some futuristic imagining equipment could give me a real-time visual impression of that event even as it unfolds. That might, perhaps, be one way for it to be true that I see the shirt's looking blue to me.

³¹ Again, facts can be equated with true propositions. It will not matter to my purposes if one wants to speak in terms of looks-facts or looks-statements.

comparative looks-facts, then demonstrative looks talk would not fundamentally add anything to his argument in terms of the latter facts. Again, that would be bad for Travis's case. By Travis's own lights, then, it would seem to be attractive to think of demonstrable looks talk as providing a somehow deeper, or anyway more lucid, view of what is in question when something looks blue to one, or when someone looks like her sister to one.³² This was the interpretative conjecture introduced at the end of section 2 above. As far as I can read the dialectical situation, it does not seem unfair for the representationalist to rely on it, and I will do so for the rest of the section.

Several of Travis's remarks suggest that, on his view, the incoherence argument, or something very much in the same spirit, could be run as well, or even better, in terms of demonstrable looks.³³ I will now consider a way of developing the argument in these terms. Suppose, then, that a shirt looks blue to me. This is to say there is a certain

³² Someone might protest that we are overlooking a third option here, viz., a 'no-priority view', on which demonstrative looks talk, on the one hand, and qualitative and comparative looks talk, on the other hand, are equally basic ways of talking (for our purposes). The 'no-priority' view might take two forms, depending on whether the two forms of talk are independent of each other or not. The idea that the forms of talk are mutually independent is evidently implausible: clearly what look something has cannot be independent of the way it looks. Such a view is moreover problematic given Travis's purposes. Our original representationalist rejoinder to the incoherence argument proposed that qualitative looks-statements – statements of the form 'x looks F' where 'F' is a qualitative adjective – are suited to ground coherent representational contents for perceptual experiences. So far we have found no good reason in Travis to think that this proposal is wrong. If demonstrable looks talk is independent of qualitative looks talk it is hard to see how one could undermine this representationalist proposal in terms of it. The other version of the no-priority view, on which the two forms of talk are closely interconnected by none prior to the other, is also problematic. For one thing, it has the consequence that the question what it is for someone to see a certain demonstrable look had by an object cannot be answered in terms of the object looking a certain way. So some other account of what the seeing of demonstrable looks involves must be offered. As I shall argue below, no account seems to be available to Travis that is both independently plausible and without a covert commitment to representational content.

³³ E.g., 'A demonstrable look, like any other visible thing, may sometimes indicate something. But that is a contingent matter' (p. 80); '[N]o demonstrable look, nor any look in our first sense, *has* any intrinsic import' (p. 81); 'One could draw the look in question [i.e. the look a pig has when it stands before me: AN]. It is, if anything, a demonstrable look; thus equally a way things would, or might, look if any of countless other things were so – for example, if there were a peccary before me' (p. 91).

demonstrable look, which the shirt now has, and which I am seeing. Now, there are countless other things that, as we say, do look, or would look, the same to me as this shirt now does. That is to say, there are countless other things that have the same demonstrable look as it.³⁴ Any one of these countless other things may or may not be of the same kind as the shirt, but they are certainly of different and mutually incompatible kinds relative to each other. As seen in a given case, this demonstrable look may indicate that one is seeing something of one kind rather than another: a blue thing rather than a white thing bluely illuminated, say. But this is contingent and not a matter of representation. Now, suppose there is something about the look seen in the original case of the shirt that makes it represent that the seen thing is blue. If this is how it is in the original case, we have no reason to deny that the demonstrable look seen when we look at the white shirt in blue light represents that the seen thing is white and bluely illuminated. Since it is the same look, or a look of the same visible type, that is present in these two situations, we find that the look determines an incoherent representational content, if it determines any.

On the face of it, this new version of the incoherence argument at least one attractive feature compared to the old version, viz. that, on the face of it, it in no way relies on the discredited idea that any qualitative looks fact can be understood in comparative terms.

Some representationalist might want to take issue with this argument at the point where it denies that indication is representation. Perhaps some sort of indication, appropriately constrained in terms of a suitable, non-circular notion of the normal visual conditions in

³⁴ One might think that these other things do not have numerically the same demonstrable look but rather demonstrable looks of the same visual type. This seems not to be the way Travis conceives of it, however. He seems to allow that different physical objects can have the same demonstrable look: 'The look in question is demonstrable, shared by anything looking suitably the same as an exemplar' (p. 74).

which the look arises, amounts to representation.³⁵ If the representationalist had to concede everything else about the argument, I believe she would have to concede too much. The view she would be saddled with would seem to be more akin to traditional indirect realism than she would want, even if it were a version of indirect realism that allowed a form of representational content to apply somehow to the direct objects of perception.³⁶

Other representationalists might be tempted simply to deny that there are any such things as demonstrable looks. Perhaps this is something one ought to deny in the end. However, there undeniably is a natural everyday use of 'look' as a count noun. We say for example that someone likes the look of their new wallpaper. The count noun corresponding to 'appear', '(an, the, ...) appearance', is also commonly used, e.g., when we say that the appearance of the couch from this side is rather different than its appearance from the other side. It would be good to avoid the consequence that all talk in these terms is simply erroneous. A better response for the representationalist is to press the question: how is demonstrable looks talk to be understood? What, quite simply, are demonstrable looks?

Now, Travis says rather little directly about what demonstrable looks are. I will not try to reconstruct what his view of that they are might be, on the basis of the various remarks he makes about them. Rather, what I aim to do is to chart what options we have when it comes to understanding demonstrable looks. I will argue that none of these options are both available to Travis and independently tenable.

³⁵ Travis briefly addresses this sort of response on p. 92.

³⁶ We shall get back to the sense in which the view is akin to indirect realism below.

5.1 Contextual paraphrase in terms of looking a certain way

Perhaps the most natural account of what demonstrable looks are is the following: what is for someone to like the look of the new wallpaper is for her to like the way the new wallpaper looks; what it for the appearance of the couch from this side to be different from its appearance from that side is for it to appear, or look, a different way from this side than it does from that side. Briefly, we naturally spell out demonstrable-looks talk in terms of things looking a certain way. The phrase ‘looking a certain way’, as used in the last few remarks, is plausibly regarded as a schematic phrase the instances of which include ‘looking blue’, ‘looking like her sister’, ‘looking happy’, ‘looking as though he had been swimming with his clothes on’, and so on. That is to say: it is plausibly regarded as a schematic phrase the instances of which are precisely phrases in terms of which qualitative (x looks F) and comparative (x looks like a G) looks-facts are specified. The natural account, then, accounts for what is in question when something has a certain look in terms of qualitative or comparative looks-facts. That is the converse of the ‘order or explanation’, or order of elucidation, we found Travis to be committed to above.³⁷

5.2 Demonstrable looks as sense-data

The order of elucidation we found Travis to be committed to is exemplified by the account Jackson (1977) offers of phenomenological looks statements (a subclass of

³⁷ It is possible that Travis, in the light of his broadly Wittgensteinian and Austinian sympathies, would squirm at the suggestion that he so much as *could* be committed to treat one type of everyday looks talk as ‘explanatorily prior’ to another, closely related type of looks talk. Perhaps he would protest that he is not so much as in the business of trying to explain anything – that he, as Wittgenstein (1942: 18) advised, self-consciously abstains from such pretensions. Nonetheless, he might still be able to agree that one way of talking of looks may allow one to ‘command a clearer view’ of things, for certain purposes, than other, related ways of talking of looks. The reference to order of elucidation can, if one wishes, be taken to mean no more than that.

what we have called qualitative looks-statements, restricted to terms for shapes, colours and spatial relations). Jackson writes,

We noted in §3 that a statement like

(22) *X* looks blue to *S*

Entails something like

(23) There looks to *S* to be something blue.

On the Sense-datum theory, the explanation for this is that (22) is equivalent to

(24) *S* immediately sees a blue sense-datum belonging to *X*. (1977: 104)

Jackson not only takes (24) to be equivalent to (22) here, but also to be analytically or explanatorily prior to it: to lay bare the ‘ontic commitments’ (1977: 107) of such phenomenal looks statements as (22), and account for some putative inferential relations among them (1977: 104-5). Travis need not, perhaps, take demonstrable looks talk to be prior to qualitative looks-statements in these two precise ways, but he is, on our reading of the dialectical situation as set out earlier in this section, committed to treat such talk as *somehow* explanatorily prior to qualitative looks-statements. Aside from treating sense-datum talk as prior to phenomenal looks-statements, there are at least three further points of analogy between Jackson’s sense-datum account and Travis’s view. First, sense-data are things subjects see. Second, the fact that we can and sometimes do see sense-data that are distinct from such physical objects as tables and chairs and cats and dogs, or parts or arrangements of such objects, accounts for how it can be that sense experience can be and sometimes is illusory or hallucinatory.³⁸ Third, while sense-data

³⁸ In fact, Jackson (1977: ch. 5) thinks that sense-data are always distinct from (parts of) physical objects, since he holds that sense-data are always coloured while physical objects never are. However, were it not for this view of his on the metaphysics of colour, it would be consistent with his other views to hold that sense-data are only sometimes distinct from physical objects. The latter claim would suffice to make the point in the text. It would also reinforce the parallel with Travis, since there is nothing in his

may indicate that things are thus and so, they do not as such represent anything as so. What Jackson affirms of sense-data here, Travis affirms of demonstrable looks.

Our question about demonstrable looks was: how are they to be understood? The suggestion that they are to be understood in terms of things looking a certain way did not fit the dialectical situation. Another answer now suggests itself in light of the last paragraph, viz., that they are sense-data. By 'sense-data' I mean non-physical particulars that bear a restricted range of qualities, are objects of visual awareness, and are such that our immediate awareness of them accounts for the phenomenal character of experience (at least for the 'sensory core' of that phenomenal character). The qualities that sense-data bear may include, as they do on Jackson's account, colour (red, blue, ...), shape (round, pointy, ...), and certain spatial relations (larger than, ...). Alternatively, they may include certain 'phenomenal analogues' of colour, shape, and spatial relations, i.e. colour' (red', blue', ...), shape' (round', ...), and spatial-relatedness' (larger-than', ...).³⁹ In any case, I take it that Travis would not want his demonstrable looks to be sense-data. Even if demonstrable looks should take the form of sense-data only in so far as experience is hallucinatory, or illusory, the resulting view

views that commits him to say that demonstrable looks are always distinct from distal physical objects, or parts of them. It might also be recorded, at this point, that Jackson, unlike most traditional sense-datum theorists, puts little or no weight on arguments from illusion or hallucination in his positive case for his sense-datum account. However, he does note as a point in favour of the view that it offers a ready account of the visual experiences implicated in hallucination (1977: 105).

³⁹ Such 'primed predicates' as red' were made famous by Peacocke (1983: chs. 1-2; 1984). Peacocke took 'red'' to describe a 'sensory property' of visual experiences (where this sensory property is a phenomenal analogue of surface redness), not a property of things experienced. But there is natural extension of the basic idea into a sense-datum framework.

would be unattractive, since there are problems of a general metaphysical sort with the very idea that we are aware in experience of non-physical particulars.⁴⁰

There are two obvious ways in which a thing seen can fail to be a non-physical particular: it can be physical, or it can be universal. So there are two obvious alternatives to the equation of demonstrable looks with sense-data: equating them with physical particulars, or equating them with universals. I will consider these in turn.

5.3 Demonstrable looks as physical particulars

The suggestion that demonstrable looks are physical particulars of some sort immediately gives rise to the question: what physical particulars? Is the look of the new wallpaper simply identical with the new wallpaper? It seems that it cannot be, for the look of the wallpaper changes from the morning to the evening, depending on the lighting, but the wallpaper doesn't (mere-Cambridge and irrelevant changes aside). Is the look of the new wallpaper simply a temporal part of it? It seems that it cannot be, for the look of the wallpaper from the kitchen differs from the look of the wallpaper from the couch, even at one moment in time. Yet the temporal stage of the wallpaper is the same as seen from both directions. Is the look of the wallpaper a complex physical object consisting of a time-slice of the wallpaper and a cone of light reflected from it onto a given position in space?⁴¹ It seems that it cannot be, for the look of the wallpaper

⁴⁰ Metaphysical problems associated with sense-datum theory are spelled out in Tye (1984), Jackson (2000), Foster (2000: 164-170), Johnston (2004: 145-156). Some, but by no means all, of these problems arise from a background commitment to token-token physicalism.

⁴¹ Brewer (2006, 2007, forthcoming) proposes an account of visual experience, in particular of visual experience in cases of illusion, which could be taken to be *broadly* analogous to this proposal. His suggestion is somewhat different in that he does not want to say that the cone of light reflected from the distal object and transmitted to a position in space is a part of the *object* of visual awareness: it is not seen, on his view. Nevertheless, such optical facts somehow individuate, or anyway are constitutive of, the conscious type of the experience. If the wallpaper isn't purple but looks purple to you, the wall's looking purple to you now partly *consists* in optical facts about the light reaching you (forthcoming: 14).

to a perceiver with normal colour vision differs from the look of the wallpaper to a perceiver with grey-scale vision, even when they see it at the same time from the same position. Yet the time-slice-of-the-wallpaper-cum-light-cone is the same for both perceivers. Is the look of the wallpaper an even more complex thing that consists of the latter plus some relevant pieces of machinery in the brain? It seems that it cannot be, but now for a different reason. Remember that demonstrable looks are supposed to be things that are seen. But we do not see pieces of machinery inside our own brain.⁴²

Brewer's proposal is intriguing, but I remain unconvinced. In addition to various problems he raises for himself, I would like to note the following. First, there is an internal problem to his view. Brewer purports to defend the 'Object View' (which he traces back to Locke and Berkeley) according to which 'the core subjective character of perceptual experience is to be given simply by citing its direct object' (forthcoming: 6). It is reasonable to read 'core' here as alluding to a 'sensory core' as opposed to higher-level features of conscious visual awareness such as, perhaps, kind properties (looking to be an X-ray machine), social properties (looking expensive), and the like. If this is how 'core' is to be understood, then the wallpaper's looking purple is certainly part of the core subjective character of your visual experience. Yet, since Brewer excludes the light-cone from the 'direct object' of visual awareness, on his own account the core subjective character in this case is *not* given *simply* by citing its direct object (since this includes, on his account, only the wallpaper). This means that, if he wants to remain true to the Object View, there is pressure on him to move towards the type of proposals set out in the text. Now, perhaps 'core' is to be understood in some other way, in which not even such a simple and basic feature of visual phenomenology as the wallpaper's looking purple is part of it. However, such a restriction on 'core' would raise the worry that the Object View is rather trivial.

Second, not all illusions need be explainable by objects or events outside the perceiver's skin. At any rate, it certainly seems possible for there to be illusions that arise merely through some abnormalities of internal processing. The arguments to follow in the text exploit this possibility. If a third person, standing beside you, has an illusion in which the wallpaper looks turquoise, and this illusion is of this internally determined kind, the wallpaper's looking turquoise to him cannot consist merely in facts about the light reaching him, since this will not distinguish his illusion from yours. To get the requisite distinction it seems we need to appeal to facts going on inside the third person's brain. As we will get back to in the text, such an appeal would be hard to square with revising Brewer's view in the manner recommended by the aforementioned problem, since things going on inside one's brain are not plausible objects of one's visual awareness.

It also leads to the following third problem. For reasons of multiple realisability, it would be unattractive to type-individuate visual experiences, even in part, with appeal to specific, physically described, events in the brain. Perhaps Brewer would say that he is not committed to such type-individuation, that he only wants to say that the wallpaper's looking purple partly consists in facts about the brain-processing. Yet if this constitutive claim is a mere implementational claim, there are reasons of a familiar sort to think that a further, higher-level characterisation of implemented fact needs to be added to capture relevant generalisations.

⁴² It is true that one need not see each part of a physical object in order to see it. I can see a cup without seeing the inside, see an apple without seeing its seeds, etc. Since the look, as construed on the

Each of the above suggestions equates the look of the wallpaper with some sort of physical object – some sort of physical continuant – even if one of a rather gerrymandered sort. Yet physical particulars also include events. Perhaps demonstrable looks are to be thought of as physical events of some kind? One might propose, for example, that the look of the wallpaper to the subject with normal colour vision (let's call him Huey) consists in the event of the wallpaper's being blue, being so-and-so illuminated, and so on. The look of the wallpaper to the subject with grey-scale vision (let's call her Grey), on the other hand, is an event that overlaps *largely* with the one just described for Huey, but differs crucially in not including the event of the wallpaper's being blue. This secures a different look for Huey and for Grey. Events are things we can see, moreover, so this proposal seems to escape all the difficulties noted for the proposals equating looks with physical objects.

Yet it encounters problems of its own. First, it is questionable whether we are entitled to say that the look of the wallpaper for Huey differs from the look of the wallpaper for Grey, on this account. Remember that the look of the wallpaper for Huey is a (composite) event which he sees. This includes the event of the wallpaper being blue. What could entitle us to say that the look-event which Grey sees does not include the event of the wallpaper being blue? So far as I can see, the only reason could be that Grey does not see it. This means that we are committed to take the following two claims to be, respectively, true and false:

last proposal, consists of several parts additionally to the pieces of machinery in the brain, someone might argue that we can see that look even if we cannot see the pieces inside the brain. This would be a bad way of rescuing the proposal. It has the consequence that the parts of the look that the subject with normal colour vision sees are precisely the parts of the look that the subject with grey-scale vision sees. Thus they precisely do not see the parts of the look in which it relevantly differs. This makes it hard to understand why the look of the wallpaper should be considered to be different for them in the first place.

(10) Huey sees the wallpaper be blue.

(11) Grey sees the wallpaper be blue.⁴³

However, reflection on the semantics of such reports as (10) and (11) suggests that (11) is true after all. Consider the following case. A certain firm has a director who expresses anger facially in an unusual way; it takes a trained eye indeed to distinguish his facial expression of anger from his plain, unaffected visage.⁴⁴ This is known to old hands at the firm, but not to new arrivals, of whom Frank is one. One day Frank comes across the director who happens to be angry in the way described. In this situation, it is certainly false that:

(12) Frank sees that the director is angry.

It is also false that

(13) Frank sees the director as being angry.

and that

(14) Frank sees the director to be angry.

and that

(15) The director looks angry to Frank.

However, it is true that

(16) Frank sees the director be angry.

To appreciate that (16) is true, consider a conversation between two old hands at the firm over whether Frank has ever seen the director angry. One of them points out that Frank saw the director in the corridor this morning. He recalls that the director was certainly angry in his characteristic way this morning in the corridor. On this basis, he may conclude that Frank has seen the director be angry.

⁴³ I am assuming that such 'naked infinitive' reports as (10) and (11) are to be analysed in terms of the seeing of events. On this analysis, the truth of (11) commits us to say that there is an event of the wallpaper being blue, and Grey sees it. See Higginbotham (1983) for a defence of such an analysis.

⁴⁴ I owe this example to Tim Williamson (in conversation).

In light of this, it is questionable why we should not say that Grey sees the wallpaper be blue. It is uncontroversial that she does not see *that* the wallpaper is blue, nor sees it *as* blue, nor *to be* blue. It is agreed that the wallpaper does not look blue to her. But the case of the angry director shows that these agreed points do not rule out that Grey sees the wallpaper be blue. Just as anger is visually hidden from Frank in the case of the director, the blue hue is visually hidden from Grey in the present case. The analogy argues in favour of saying that Grey does see the wallpaper be blue.

Another problem for the idea that demonstrable looks are physical events arises over illusions.⁴⁵ Let's say that Frida and Frodo are looking at the same display of red dots on a white background, in identical lighting, from the same point of view. (It is best to imagine that Frida and Frodo inhabit different possible worlds.) For some internal reason or another, the dots look orange to Frida while they look purple to Frodo. With respect to the dots, there is a difference in the colour phenomenology of their visual experiences. Talking in terms of demonstrable looks, what's going on here is that the dots have a certain look for Frida which differs from the look they have for Frodo. According to the current idea, these looks are physical events: distinct physical events, indeed, as they are distinct looks. Now, exactly the same things may be going on physically outside of Frida's and Frodo's brains; if we pick an event from Frida's environment to be the look of the dots for her, that event happens in Frodo's environment too, and vice versa. The challenge facing us now is this: how are we to find such an event that *only* one of them sees? If they both see it, we will not be able to

⁴⁵ Note that the Huey/Grey case ought not be classified as one involving illusion. We ought not say that Grey illusorily experiences a non-grey wall as grey. We ought to say that she experiences the wall as having a certain brightness, but that she doesn't experience it as having a certain hue. The wall may well have the brightness it looks to Grey to have.

explain why it constitutes the look of the dots for the one but not the other. The challenge would not be met by saying there is an event of the dots being purple that only Frida sees. That suggestion would not work even if, pace our conclusion in the last paragraph, Frodo would not have seen that event, had it taken place. The suggestion does not work because there is no event of the dots being purple. It seems that we can develop the further details of the case of Frida and Frodo in such a manner that if there is any physical event in the environment the one sees, the other sees it too. To put the same point differently: it does not seem that the slight difference in illusory colour phenomenology between Frida and Frodo entails that there *must* be some physical event around them at most one of them can see.⁴⁶ We find here the same drift as in the object case towards moving into Frida's and Frodo's brains to find things present in the one case but not in the other. To be sure there are physical events going on in Frida's brains that don't go on in Frodo's. But these are non-starters for being the look of the dots for either of them, since we do not see physical happening inside our own brain.

These considerations may not exhaust all the ways in which demonstrable looks could be equated with physical object or events. I take them to suggest, though, that looks are not physical objects or events of any kind. Physical objects and events may not, in turn, exhaust physical particulars. Perhaps there are physical particulars of some other more exotic ontological sort.⁴⁷ Still, I take what we have said to make it reasonable to conclude that demonstrable looks are not physical particulars of any sort.

⁴⁶ Even if I am wrong about this, the event in question may be a very bad candidate to the look of the dots for one of them

⁴⁷ The principal alternative not considered above are physical abstract particulars, i.e. physical tropes. Interestingly, the philosophers who have promoted the importance of tropes to metaphysics have often been keen to promote tropes as the object of sensory awareness; D. C. Williams (1953: 123) and K. Campbell (1981: 130) are two examples.

5.4 Demonstrable looks as universals

Travis, we argued, would not want his demonstrable looks to be sense-data, i.e. non-physical particulars. In the last subsection, we excluded the obvious alternative that demonstrable looks are physical particulars. This leaves us with the remaining obvious alternative: that they are universals.

Some may think that this proposal is ruled out immediately by the requirement that demonstrable looks are things we can and do see. For does it really make sense literally to see universals? It is true that it makes perfect sense to say, in natural language, such things as ‘I see the colour of his tie’. However, as G. J. Warnock (1954-55) argued, such talk is plausibly to be understood in terms of propositional ‘seeing that P’ talk. If I see the colour of his tie, I see that his tie is red, or that it is blue, or that it is ..., where some colour fills in for the dots. We are here interested in what one might call object seeing, as opposed to propositional seeing. Other talk in natural language suggestive of seeing properties, such as ‘seeing red’, is typically not literally intended.

Now, the idea that demonstrable looks are physical tropes would, at least on many views of tropes, confront the same difficulty over illusion as the physical event view was seen to. If tropes simply *are* events, as they are on some views, it would confront the problem of the Huey/Grey case too.

More generally, I would say the following. Tropes may be regarded either as metaphysically constructed or as metaphysically basic items. If they are metaphysically constructed, they are plausibly constructed from objects and properties considered as repeatables (as universals in a neutral sense of the word). It seems fair to say that most of the philosophers who are happy to allow for the existence of tropes, regard them as metaphysical constructions in this way. If this is how we look at tropes, then the idea that we can see a trope is plausibly explanatorily dependent on the idea that we can see properties. For example, the idea that we can see the trope that is this banana’s yellowness is plausibly explanatorily dependent on the idea that we can see the property of being yellow. We shall turn to the latter idea presently.

The alternative option is to regard tropes as metaphysically basic. This idea takes us into much more controversial metaphysical terrain. I am myself not persuaded that there are such entities. Moreover, I believe it would be advantageous for a philosophical account of perception to be neutral on the deep metaphysical question whether there are basic tropes. This makes a theory that is not so committed more attractive than one that is, at least other things being equal. This is obviously no knock-down objection, but does, in my book, motivate the search for a theory with less contentious metaphysical commitments.

Notwithstanding such doubts, interesting proposals have recently been made by Foster (2000) and Johnston (2004) to the effect that visual awareness of universals is central to visual consciousness. Neither Foster nor Johnston wants to say that we ‘see’ universals. We should not put too much weight on this, however. Many sense-datum theorists (Jackson is an exception) were similarly reluctant to say that we ‘see’ sense-data, preferring to say, for example, that we ‘sense’ (Broad 1923), are ‘acquainted with’ (Russell 1912), or are ‘presented with’ (Moore 1918-19) sense-data. In spite of this terminological point, sense-datum theorists took us to have the same kind of relation with sense-data as the person who explicitly or implicitly subscribes to direct realism takes herself to have to physical objects in seeing them. Similarly, Johnston, at least, takes it that we have the same kind of relation to universals in being visually aware of them as the direct realist (rightly, in Johnston’s view) takes herself to have to physical objects in seeing them. If Johnston, or the sense-datum theorists, are right about this, it seems largely a terminological question whether we shall call this relation ‘seeing’, or use some more neutral, quasi-technical term such as ‘being visually aware of’.

Foster’s and Johnston’s recent proposals differ in some interesting ways. Johnston’s strikes me as the clearest, and the most promising, so I will treat his view as the paradigm of an account where direct, visual awareness of universals is central to visual consciousness.⁴⁸ I will start by setting out an outline of his view. I will then argue that,

⁴⁸ Foster (2000: 186-196) argues that the universals that, on his view, are the immediate objects of sensory awareness are nevertheless concrete. They are thus the mirror image of tropes or abstract particulars. In some respects, his conception of concrete universals seems analogous to Armstrong’s (1978) immanent universals. One thing that remains obscure to me, however, is what the instances of Foster’s concrete universals are, or even if they are supposed to have any instances at all. If they don’t, that would wreck much of the analogy with Armstrongian universals, which exist only in their instances. The analogies he uses with temporal parts seem to suggest that the instances of his concrete universals are, in effect, in sense-data. Yet one certainly gets the impression that Foster does not want to be

in relevant respects, his view is not available to Travis, and is not an alternative to a version of representationalism, in the weak sense in which the latter term is used here.

Consider an entirely normal perceptual situation, where I am looking at my desk. On Johnston's account, I am in this situation in two equally direct, tightly interlinked states of visual awareness.⁴⁹ On the one hand, I am aware of my desk, the screen on top of it, and other familiar physical objects lying around. This act of awareness is not one I could be in if I were hallucinating. On the other hand, I am aware of what Johnston calls a 'sensible profile'. A sensible profile, here, is a complex, structural property that captures, or simply is, the way the scene before my eyes currently is, as far as it is presented in my visual experience. For example, if a paper note attached to the screen looks yellow, the property of being yellow is a part of this complex, structural property that is the sensible profile. It is a part of the way the scene before my eyes currently is, as far as things currently look to me, that there is a yellow note attached on such-and-such a location relative to me. Since the current visual situation is a normal one, we will expect that the sensible profile very largely overlaps with the way the scene before the eyes (the arrangement of the desk, the screen, etc.) in fact is. If my visual experience were perfectly veridical, the sensible profile would be a proper part of (would be entailed by, but not entail) the way the scene before my eyes actually is. This second act of awareness, my awareness of the sensible profile, is one I could be in even if I were suffering from a complete 'philosophical' hallucination. Indeed, on Johnston's view, I *would* be in this act of awareness, if I had been suffering from a complete hallucination

committed to the existence of sense-data (see, e.g., 2000:164--170). Another complicating feature is Foster's insistence that the relevant concrete universals exist in the mind – that they somehow essentially belong to an essentially mental sphere. Again, I find this rather hard to comprehend. The difficulties I have with this conception are, perhaps, associated of the fact that Foster wants his concrete universals to serve as the launching pad for an idealistic metaphysics about the physical world.

⁴⁹ We shall return below to the question how they may be linked.

and had been intrinsically just as I am. The reason why I could in this state of awareness is that the sensible profile is a universal, not dependent for its existence on any one of its instances; indeed, not dependent for its existence on its having any instances at all.⁵⁰

Johnston's account of visual consciousness is rich and detailed, and the above sketch does not begin to do justice to its many interesting features. However, it should suffice to make it relatively clear why it would not be attractive to Travis to equate seeing demonstrable looks with being visually aware of sensible profiles. Johnson's introduction of the notion of a sensible profile, like ours above, makes serious use of the notion of veridicality as applied to visual experience:

Consider the sensed field or scene before your eyes. Now attend to the relational and qualitative structure that is visibly instantiated there in the scene. It consists of just the properties and relations of which you are visually aware, when you are seeing the scene. It is a scene type or *sensible profile*, a complex, partly qualitative and partly relational property, which exhausts the *way* the particular scene before your eyes is if your present experience is veridical. (Johnston 2004: 134)

Travis, of course, forswears any serious use of the idea of veridicality or veridicality-conditions for perceptual experience. Perhaps a defender of Travis initially sympathetic to Johnston's view would suggest Johnston is using the notion of veridicality here in a

⁵⁰ The vast majority of sensory profiles will have at most one instance, since exact, detailed qualitative and relational duplication of the scenes before our eyes, as we see them, presumably is a rarity. Moreover, since sensory experience presumably is somewhat illusory in one respect or another fairly regularly, a considerable number of sensory profiles will have no instance. The possible lack of an instance will apply not only to the highly structural universals that typically constitute sensible profiles, but also to atomic universals apt to be parts thereof, such as the universal supersaturated red. The Platonist assumption that universals do not depend for their existence on having instances is thus crucial to Johnston's view. Another controversial ontological commitment is that there are structural universals, such as the universal of having two hydrogen atoms and one oxygen atom so-and-so combined. See Johnston (2004: 146-8) for discussion of these ontological issues.

pedagogical role only: it allows his readers to get a fix on the notion of sensible profile he wants, but he can, as it were, kick away the notion of veridicality once this fix is obtained. A theorist sympathetic to both Travis and Johnston might take further solace from the latter's insistence that 'experience is neither propositional nor qualia-involving' (2004: 176 nt. 2) and his tendency talk of 'Intentionalist views' in a way suggesting that they, as a class, contrast with his own (see e.g. 2004:114-115, 176 nts. 2, 4).

This would be a mistake. However Johnston's notion of a sensible profile gets introduced, a robust notion of a veridicality condition can be defined once the notion is in place. We can say that an act of visual awareness of a sensible profile is veridical just in case the sensible profile is instantiated by the scene around the subject.⁵¹ Correlatively, a notion of representational content can non-arbitrarily be applied to experience. On a relatively generic, but also relatively standard, use, the representational content of a given perceptual experience can be thought of as something characteristically associated with the experience which determines its conditions of veridicality. In this sense, a sensible profile counts precisely as a representational content.

Our theorist who is sympathetic to both Travis and Johnston cannot get around this by suggesting that it is, perhaps, a relatively peripheral aspect of the latter's view that sensible profiles are things which it makes sense to think of as possibly instantiated by

⁵¹ Derivatively, we could say that a visual experience is veridical just in case the visual awareness of a sensory profile that partly or wholly constitutes it is veridical.

the physical scene around the perceiver.⁵² This aspect is not peripheral, it is absolutely central to Johnston's conception of how the two acts of direct, sensory awareness in normal perception are interlinked. The act of visual awareness, in the case above, constituting my seeing the desk, the computer screen, and so on, is, on Johnston's view, an act of awareness of an instantiation of the sensible profile that *the other* act of awareness involved here is an awareness of.

A way of further pressing the point that Johnston is not a friend of Travis is by way of comparing the former to someone who certainly shouldn't count as a friend of the latter, viz., Michael Tye. On Tye's (2000, 2006a, 2006b) view, visual experience has representational content taking the form of an existentially quantified, structured, neo-Russellian proposition. Like Johnston's sensible profiles, Tye's neo-Russellian contents contain properties instantiable by normal physical objects as atomic parts. Like the sensible profiles, the neo-Russellian contents are built out of these basic parts by various logical or quasi-logical operations. Like Tye's existentially quantified contents, Johnston's sensible profile turn out also to be implicitly existentially quantified:

The sensible profile or *way* the scene is involves more than the layout. For example, it includes the further condition that the relational layout be filled in with some particulars or other that have such and such qualities. (2004: 135)

⁵² Foster (2000) seems to want to think of the concrete universals that are the immediate objects of awareness on his view as simply not the right kind of things to be instantiated by mind-independent objects. However, he also seems to want *not* to posit sense-data as objects instantiating them, nor to portray them as 'transcendent', Platonic universals. I don't see how he manages to square these desiderata. Perhaps his idealist metaphysics about the physical world allow him to do this in some way. Yet if idealism is required to make his conception of the immediate objects of awareness coherent, adoption of it comes at a considerable price.

These points suggest that Johnston's sensible profiles can be correlated one-one with Tye's propositions. I cannot see that Johnston says anything about sensible profiles that would prevent such a correlation. True, whereas Johnston stresses that sensible profiles are objects of visual awareness, Tye tends not to talk in that way of his contents, saying rather that representational contents are things perceptual experiences have. However, there does not seem to be anything in principle that would prevent him from talking of his neo-Russellian contents as objects of awareness.⁵³ The impression one gets, then, is that a core portion of Johnston's account, devoted to the existence and fundamental nature of perceptual awareness of sensible profiles, and a core portion of Tye's account, devoted to the existence and fundamental nature of perceptual experience having representational content, can be regarded as effective notational variants of each other. This not to deny, of course, that their further views on perceptual consciousness differ in several interesting respects.

The conclusion we may draw is that Travis, or anyone else who wants to deny representationalism even in our weak sense, ought not identify the seeing of demonstrable looks with the visual awareness of universals, if such awareness is understood in anything like Johnston's way. True, there are alternative ways of understanding sensory awareness of universals, e.g. Foster's. Yet Foster's conception of the 'concrete universals' that, on his view, are the immediate objects of sensory awareness is hard to comprehend, and may require accepting his idealist metaphysics.⁵⁴ I am not sure how the other alternative conceptions would go, but see no grounds to think that they would offer the anti-representationalist better candidates for the role of

⁵³ Such a way of talking may indeed fit in rather well with some of Tye's arguments for his view, such as the weight he puts on transparency considerations (2000: 45-64).

⁵⁴ For some critical observations on Foster's view, see fn. 48 and 52 above.

demonstrable looks. The reasonable conclusion is that universals are not well suited to play the role the anti-representationalist wants demonstrable looks to play.

We have now surveyed a range of ways of understanding Travis's notion of a demonstrable look. The most natural account of demonstrable looks says that having a certain demonstrable look is a matter of looking a certain way. It accounts for demonstrable looks-facts in terms of qualitative (x looks F) and comparative (x looks like a G) looks-facts. If this is how we ought to conceive of things, Travis's claims and arguments in terms of demonstrable looks don't add anything to the same claims and arguments as reformulated in terms of qualitative or comparative looks-facts. In particular, the version of the incoherence argument formulated in terms of demonstrable looks wouldn't add anything to the incoherence argument formulated in terms of qualitative and comparative looks-facts. Since the latter did not work, this would be an unfortunate conclusion for Travis. In light of this dialectical situation, he seems to be committed not to endorse what we have been calling the most natural account.

Three alternative accounts of what demonstrable looks are are, respectively, that they are sense-data, physical particulars, and universals. The first option seems to fit in with many of the features Travis attributes to demonstrable looks, but is unattractive on metaphysical grounds (if no others). The second option does not allow us a sufficiently fine-grained account of looks, or forces us to say, implausibly, that one sees things inside one's own brain. We have found no reason to think that the third option is false. Indeed, the idea that demonstrable looks are universals strikes me as pretty plausible. The problem is that this view, on its clearest and most promising development, is simply a version of representationalism; it is not an option for an anti-representationalist

such as Travis. Now, if demonstrable looks, on the anti-representationalists view, are not capable of being paraphrased away as ways things look, and can be neither non-physical particulars nor physical particulars nor universals, the reasonable conclusion is that they are nothing. The reasonable conclusion is that they cannot do the job Travis wants them to do.

6. Seeing and representational content

Some of Travis's arguments against representationalism are independent of whether or not representationalism is committed to looks-indexing, or anyway purport to be. Central among these is an argument to the effect that the representationalist is unable to give any plausible account of the relationship between one's having an experience with representational content and one's seeing physical objects (pp. 87-91).⁵⁵ Travis distinguishes two ways in which the representationalist may conceive of this relationship, calling them the 'one-source' and 'two-source' models, and argues that none works. He describes the two models so:

On the one, I have two separate sources of information. I am aware of two different things: first, the pig; second, being represented to. Awareness of neither carries with it, ipso facto, awareness of the other. I see the pig; it is otherwise intimated to me, redundantly, that there is one. On the other, there is but one source. I am represented to, and cognizant of that. To be aware of all that I thus am just is, where I do see a pig before me, my doing that. It may constitute that awareness of a pig before me that seeing one (to be there) is. (p. 88)

⁵⁵ Travis introduces the argument by asking: '[I]n cases where we do see things, how does our being represented to relate to our seeing what we do?' (p. 85) He clearly purports to argue that no sound answer can be given to this question.

On the two-source model, then, whenever I see a pig, I am, on the one hand, in a state of seeing the pig, and, on the other, in a state of having a visual experience with a certain representational content, perhaps to the effect that there is pig in front of me. The former state is not one I could be in were there no pig before me, but the latter is. Travis takes it that each state independently gives me a reason to take it that there is a pig before me, though a reason of a different kind; he suggests this is problematic. He also seems to assume that, if seeing a pig involved this duality of states, the duality would somehow be phenomenologically manifest to the perceiver. The phenomenology would be like that of having ‘a relative at the movies, reciting what you have just seen’ (p. 88). It is not hard to disagree with Travis that, if this is something the two-source picture implies, it is obviously absurd.

On the one-source model, there is just one properly visual, mental, experiential state involved in one’s seeing a pig. This experiential state has a certain representational content, and is one the subject could have been in even were there no pig before her senses. What makes it so that the subject counts as seeing a pig, in a given case, is that, in addition to being in the relevant state, certain suitably pig-involving surrounding conditions are satisfied in the relevant case. Travis argues that this account fails because ‘awareness of something else plus satisfaction of surrounding conditions cannot add up to that awareness of porcine presence which we have in seeing one (to be there) – in my terms, unmediated awareness’ (p. 89). In unmediated awareness, Travis writes, ‘one’s entitlement to take it that X is hostage to no more than some form of awareness of X itself’ (p. 65). The one-source model falsely portrays seeing a pig as a case of mediated awareness because it makes one’s entitlement to take it that there is a pig present hostage not simply to a state of awareness constituting, on its own, awareness of the pig,

but rather to a the visual experience as of there being a pig before one, and the surrounding conditions being of the right kind.

Both of these arguments raise several questions. However, as I am inclined to agree with Travis that the one-source model is problematic, I will concentrate on his remarks about the two-source model. It quite unclear why the duality this model posits should be manifest to introspection. To see this, it is useful to compare with two other cases where it is independently plausible to assume the analogue of a ‘two-source’ model. Consider, first, singular, perceptual-demonstrative thought. Several philosophers have found it plausible that the perceptual-demonstrative thought ‘That bottle is empty’, unlike the quantified thought ‘Some empty bottle is before me’, is not one I could have entertained if a bottle were not visibly before me.⁵⁶ The quantified thought is one I will have both in the actual case and in the case where a Cartesian demon is deceiving me; the demonstrative thought is one I will have in the former only. As compared with the delusive case, then, there is a certain duality of states involved in the non-delusive case. But we are not to expect that it will be obvious to introspection that there is something extra in the non-delusive case, or that something is missing from the delusive case.

The second analogy concerns knowledge. Williamson has recently argued that knowledge is a mental state that entails, but is not entailed by, a combination of belief, justification, plus the satisfaction of certain environmental conditions.⁵⁷ In a normal case, where I know by looking that it is raining outside, I believe that it is raining outside *and* I am in a non-identical mental state of knowing that it is raining outside. In a subjectively indiscriminable case where I am deceived by a Cartesian demon, I am

⁵⁶ See Evans (1982) and McDowell (1979, 1984) for a defence.

⁵⁷ See Williamson (2000).

only in the former state. As compared with the delusive case, then, there is certain duality to the non-delusive situation. But we are not to expect that it will be obvious to introspection that there is something extra in the non-delusive case, missing from the delusive case.

If either of these accounts are plausible, and I believe they are, it raises the question why the duality in the non-delusive case to which the two-source model is committed should be any different. It is useful to emphasize, at this point, that the specification of the two-source model leaves fairly open how the object-dependent visual experience of the pig, entailing the pig's presence, is to be understood more precisely. One option for the representationalist would be to take this state to have a singular, object-dependent content; for example, a content expressible (in part) as '*That* thing is a pig and fat'. Assuming, for the sake of argument, that the object-independent visual experience implicated in seeing the pig has the content that *there is* something that is a pig and fat and before me, the relation between the object-dependent and object-independent states is of course highly analogous to that between the perceptual-demonstrative and existential thoughts described above.⁵⁸ There would be no more reason to expect the presence of the object-dependent state in the non-delusive case, or its absence in the delusive case to be obvious to introspection. There would certainly not be any reason to think that object-independent visual experience should manifest itself, phenomenologically, as a commentary on what one sees.

⁵⁸ It is disputed whether kind-properties ever enter the content of perception. So a more neutral specification of the respective contents might be that the object-dependent visual experience (the seeing) has the content that *that* thing is round and blue and ..., while the non-object-dependent experience has the content that there is something round and blue and See Siegel (2006) for more on the issue of what kind of properties are represented in perception.

Nor does this style of account need to imply that the perceiver has ‘, twice over, reason to take it that there is a pig, each reason different in kind, each independent of the other.’ (p. 88) The two-model view may be combined with an externalist view of reasons according to which it is only the object-dependent experience that gives reason to think that there is a pig before one. This is not ad hoc; there are parallel grounds for thinking that it is only my knowledge that it is raining outside, in the case described above, that gives me reason to take it that the streets are wet, not my belief that it is raining outside, entailed by my knowledge.⁵⁹ This externalist view of reasons implies that in the delusive case, where the Cartesian demon is deceiving me, I have no reason for my belief that there is a pig before me. Epistemic internalist generally find this unacceptable, for they take to imply that I would be somehow epistemically blameworthy in believing that there is a pig before me, in the deceptive case, which seems wrong. However, people may be blameless in doing something they had no reason to do, so this internalist objection is not decisive.⁶⁰

Even if this is wrong, and the externalist view of reasons must be given up, it does not mean that the two-source model is committed to say that the two visually experiential states implicated in seeing the pig give two different, independent reasons for taking there to be a pig. For the two-source model may equally be combined with an internalist view of reasons according to which it is only the object-independent experience that gives reasons to think that there is a pig before one. Again, this is not ad hoc, for the internalist takes reason-giving states in general to be subject to the requirement that the subject can recognise ‘from within’, in a manner that can not be upset by whether she is

⁵⁹ If a state of mind gives reason for thinking that so-and-so only if it provides evidence for thinking that so-and-so, and Williamson (2000: ch 9) is right that only knowledge provides evidence, only my knowing that it is raining outside gives me reason to think that the streets are wet.

⁶⁰ See Williamson (2007) for further discussion and defence of this point.

in a sceptical scenario or not, what reasons she has. The object-dependent visual experience does not meet this requirement. On either of two major views of reasons, then, the two-source model does not have the consequence for reasons the Travis ascribes to it.

What if both the purely externalist and purely internalist views gestured at above are wrong? Then the two-source theorist may well want to say that two types of reason-giving visual states are implicated in seeing a pig, giving reasons of different types or in different ways. Yet far from being objectionable, this just what one would expect on general epistemological grounds if both pure externalism and pure internalism about reasons is wrong.

The two-source model of seeing need not, then, have either of the objectionable features Travis ascribes to it. Alternatively, if the putatively two-source accounts we have outlined above do not count as such on Travis's intended use of 'two-source model', his two models fail to capture a promising conception of how visual representation may figure in the seeing of pigs. In any case, his argument that the representationalist is unable to give a plausible account of the relationship between one's having an experience with representational content and one's seeing physical objects fails.

7. Conclusion

The overall message of this paper is that Travis's case against representationalism ultimately fails. This is of course not to say that his case is of little importance; on the contrary, his arguments raise several deep and critically important questions for representationalism, many of which deserve fuller treatment than we have been able to provide here. Another thing that is not part of the message of this paper is that

representationalism has been shown true. However successful we have been in rebutting Travis's objections to the view, we have not vindicated its truth, or purported to. All the same, I believe that going over the ground we have covered will prove useful to a positive case for representationalism, by drawing attention to some points in the dialectical landscape which would profit from further investigation. One such point is surely the semantics and metaphysics of looks- and appears-talk, and the case for the claim we called Non-comparative Grounding. This is a good time to take up the thread from Chisholm (1957) and Jackson (1977) seminal work on these issues.

Conceptualism, Recognitional Concepts, and Indiscriminability

Abstract

Some conceptualists are committed to the idea that subjects have recognitional capacities for the specific features they experience things as having. I argue that subjects do not have such capacities, not because of limitations of our sensory memory capacities (as several nonconceptualists have argued recently), but because of the non-transitivity of phenomenal indiscriminability, or matching.

0. Introduction

According to weak representationalism, perceptual experiences are analogous to beliefs and desires in at least the following respect: just as beliefs and desires have representational contents that determine, respectively, the truth-conditions of the beliefs and the fulfilment-conditions of the desires, so perceptual experiences have representational contents that determine the veridicality-conditions of the experience.¹ The many theorists who agree on this much disagree, however, on whether perceptual

¹ In a recent survey, Chalmers (2004) remarks that weak representationalism is 'rarely denied'. While this is certainly a largely correct description of contemporary opinion, some interesting objections to the view has recently been advanced, by among others John Campbell (2002) and Charles Travis (2004). I defend weak representationalism against Travis's objections in essay 1 above, and against Campbell's in Nes (2006: ch. 6).

experiences are analogous to beliefs and desires in any one of several other candidate respects.

One of these candidate respects is whether perceptual experiences have their content in a conceptualised way. It is widely assumed that beliefs and desires are conceptualised in the following sense. Someone can have a belief that gravity acts on a falling dinner plate, for example, only if she has a concept of gravity. Someone can have a desire that she gets antibiotics soon, only if she has a concept of antibiotics. More generally, we take beliefs and desires to be subject to the constraint, roughly, that someone can believe, or desire, that ...F... only if she has a concept of F/F-ness/....² *Conceptualists* affirm, while *nonconceptualists* deny, that this constraint generalises to perceptual experience, i.e. that someone can have a perceptual experience with the content that ...F... only if she has a concept of F.³ Speaks (2005) has recently argued, I believe rightly, that conceptualism should be the default assumption for weak representationalists, since it extends to perceptual experience a constraint that holds across a wide range of intentional states.

Yet conceptualism is subject to a putative counterexample, noted independently by David Kelley (1980) and Gareth Evans (1982). Perceptual experience is extremely

² The slash and the dots are there to allow that different instances of this schema may have to be formulated slightly differently depending on the semantic type of 'F'. See the Introduction, section 3.1, for more on this problem. I ignore this complication in what follows.

³ Influential defences of conceptualism include Armstrong (1968), McDowell (1994), and Brewer (1999). Influential defences of nonconceptualism include Evans (1982), Peacocke (1992, 2001), and Tye (1995, 2000). Byrne (2005) (developing suggestions in Stalnaker (1998) and Heck (2000)) makes explicit a distinction between *state* and *content* (non)conceptualism. In this paper, only state conceptualism is directly in question. The relation between state and content conceptualism is discussed in the Introduction.

specific, or fine-grained, in its content. For example, we are capable of sensibly discriminating over ten million shades of colour, according to some estimates.⁴ Suppose that red₁₂₃ is one of these highly specific shades of colour we may experience something as having. Conceptualism has the consequence that to experience something as red₁₂₃ a subject needs to possess a concept of red₁₂₃. This might seem highly implausible: ‘Do we really,’ as Evans rhetorically asked, ‘understand the proposal that we have as many colour concepts as there are shades of colour we can sensibly discriminate?’ (1982: 229)

A conceptualist response to this problem worthy of serious attention must do more than assert that we have do have such concepts. It must explain what entitles us to think that we do. John McDowell has offered a response of this kind.

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 exactly as fine-grained as the experience, by uttering a phrase like ‘that shade’, in which the demonstrative exploits the presence of the sample.... What is in play here is a recognitional capacity, possibly quite short lived, that sets in with the experience. (1994: 56-7)

We may distinguish two strands in McDowell’s response here. The first is that, on experiencing an object as red₁₂₃, we have a concept *D* of that shade, to which we might give expression by means of phrase like ‘that shade’. We may think of this as the ‘demonstrative strand’ of his proposal, according to which we need to possess perceptual-demonstrative concepts of the features we experience things as having. I am using ‘perceptual-demonstrative concepts’ in a broad sense here to denote concepts that

⁴ See Tye (1995).

presume a perceptual relation to the referent and which are naturally or canonically expressed in natural language by means of demonstratives.⁵ In this paper, I will grant this strand of his proposal, for the sake of argument. The second strand is that the concept *D* is associated with a recognitional capacity for red₁₂₃ things. The idea is that our claim to think of the subject as possessing a bona fide *concept*, in possessing *D*, presumes she has such a recognitional capacity. The idea seems to be, that is to say, that possessing *D* presumes a capacity for recognising some other red₁₂₃ things as also falling under *D* when appropriately presented with them. The appropriate way of presentation in question here is, plausibly, the way realised just when I see the object in question and it visually presents the shade red₁₂₃ to me. To say that I can recognise a red₁₂₃ object so presented as falling under *D* is to say that I am in a position to know, on the basis of a perceptual experience of this kind, that *that* thing is *D* too.⁶

Now, an influential objection to the second, ‘recognitional strand’ of McDowell’s proposal appeals to certain empirical limitations on our sensory memory capacities. Research in experimental psychology strongly suggests that subjects are capable of discriminating much finer shades of colour at a time than they are capable reliably of re-identifying across time. This is so even if the objects to be classified as either the same or not the same as the target are presented subsequent to just the briefest interval after the target has passed from view. If this is so, McDowell might seem to face the

⁵ For a contrasting, stricter use of ‘demonstrative concepts’, see Papineau (2007). I will let the parenthetical ‘perceptual-’ in front of ‘demonstrative’ be tacit henceforth.

⁶ Two comments: (i) I am presuming that recognising *x* as *F* requires knowing that *x* is *F*. (ii) Our construal of the relevant appropriate way for red₁₂₃ things to be presented entails that it is not enough, to count as having a recognitional capacity of the relevant type, that one can know, on the basis of testimony, that that thing is *D* too, or know this by inference from knowledge of the physicalistic supervenience base for colours, say. McDowell clearly seems to have a *perceptual*-recognitional capacity in mind.

following dilemma: *either* D is not a concept of the most specific shade we can sensibly discriminate the object as having at a time, and so, *pace* McDowell, is not ‘exactly as fine-grained as the experience’, *or* if it is such a concept, D is not associated with a recognitional capacity, again *pace* McDowell.⁷

However, it seems the conceptualist has a reply to this objection, in as much as he might deny that the recognitional capacity associated with D *needs* to outlast the time t when the original red₁₂₃ object, let’s call it o , is presented. It would be coherent to hold that that capacity is retained only as long as o is in view, and only as long as one’s experience of o , colour-wise, remains unchanged. The claim that one has a recognitional capacity (only) then is not empty since it supports counterfactuals to the effect that if another red₁₂₃-looking, red₁₂₃ object were presented next to o at t , the subject would be able knowledgeably to judge of it *that thing is D too*. Given this weaker ‘merely cross-modal’, as opposed to ‘cross-temporal’, constraint on having a recognitional capacity, the facts about limitations on memory to which several critics have appealed do not refute the idea that D comes with a recognitional capacity.

Nonetheless, in this paper what I will argue is that, even given this weaker necessary condition for having a recognitional capacity, the idea that D comes with a recognitional capacity is untenable. The argument rests on the widely presumed non-transitivity of phenomenal indiscriminability or ‘matching’.⁸ If, and only if, matching is non-

⁷ Raffman (1995), Peacocke (2001) and Kelly (2001) press this kind of objection to conceptualism. See Raffman (1995) and Kelly (2001) in particular for references to the original experimental research.

⁸ The notion of phenomenal indiscriminability, or of matching (I use these interchangeably), is understood in a variety of ways in the literature. Graff (2001) identifies matching with the relation of looking the same, where the sense of ‘looks’ in play is supposed to be one with no explicit

transitive, one can construct so-called phenomenal sorites series. A (visual) phenomenal sorites series is a series of objects x_1, \dots, x_n such that each object, x_i , (visually) matches the next, x_{i+1} , in colour, say, although the first, x_1 , does not (visually) match in colour the last, x_n .⁹ We shall see how this motivates a certain epistemic challenge to McDowell.

McDowell writes,

In the presence of the original sample, a subject who has the concept of a shade is enabled to classify items as possessing the relevant shade or not, by direct inspection for colour match. (1994: 172)

Now, let's grant, at least for the sake of argument, that one, 'by direct inspection for colour match', is in a position to know whether or not a given object x matches the original sample o in colour. Yet, since matching is non-transitive, this would appear to be compatible with x having a different visually presented shade from the visually presented shade of o ; it would appear to be compatible with the shade x looks being distinct from the shade o looks. So even taking for granted that the object has the shade

epistemological implications. Williamson (1990) by contrast defines an overtly epistemological notion of indiscriminability, according to which x and y match (relative to a subject, a time, a source of information, or other parameters) iff it is not possible to know (within those parameters) that x and y are different. Hardin (1988) characterises yet another, statistical or response-dependent notion of matching, where whether or not a target matches a sample depends an observer's ability to do better than random at distinguishing the target after sufficiently many trials. However, a number of writers who discuss the transitivity or otherwise of matching or phenomenal indiscriminability do not discuss what notion they are invoking. This in particular applies to the conceptualists writers who are my targets in this paper, viz. McDowell and Brewer. As far as I can tell, it does not matter to the argument of this paper whether we adopt a non-epistemic, phenomenal notion of matching like Graff's, or an epistemic one like Williamson's. However, if pressed to choose, the former, phenomenal notion should be taken as canonical for my purposes here.

⁹ Classic discussions of the presumed non-transitivity include Goodman (1951), Dummett (1975) and Wright (1975). Although widely agreed, non-transitivity is not entirely uncontroversial; two recent critics are Raffman (2000) and Graff (2001).

it looks or visually presents, it is not clear why our subject should be in a position knowledgeably to classify x as having the most fine-grained shade o visibly has, since it seems that, for all she knows, x doesn't. This is the prima facie epistemic challenge to McDowell.

I now turn to consider McDowell's (1994) and Brewer's (1999) positive treatment of phenomenal sorites series.¹⁰ I examine their treatment since it is the only conceptualist treatment of this issue of which I am aware, and since it is highly interesting and prima facie promising.

1. McDowell and Brewer on phenomenal sorites series

I have talked in passing about there being a certain 'original object' for a demonstrative concept such as D . The idea is that the demonstrative concept is made available by a particular object of perception: even though the concept is of a property, it is made available by a particular object which visibly exemplifies that property. The key idea in McDowell's and Brewer's treatment of sorites series is that this original object has not only a causal and enabling tie to the demonstrative concept, but also a more intimate semantic tie of some kind with it. McDowell writes,

An utterance of 'that shade' depends for its meaning on the identity of a sample shade. We might lay down the rule that *something counts as having that shade just in case it is indiscriminable in colour from the indicated [i.e. the original] sample*. (1994: 170, my emphasis)

McDowell refers to the italicised rule here as 'the first rule' and I will use that term:

¹⁰ For present purposes, any differences between their treatments do not matter.

The First Rule An object x has D_o iff x is indiscriminable in colour from o .¹¹

Remember, the demonstrative concept D_o purports to present a property, a universal. The original sample, o , has that property, all right, but so far we have been given no reason for thinking that it should be privileged in determining what particulars falls under the concept D . Suppose we judge, in accordance with the First Rule, that a certain object x_1 , non-identical with o but indiscriminable from it in colour, also has the demonstrated shade D_o . Why should not x_1 be as good a candidate as the original o for the role of determining the extension of the shade concept? The reason, McDowell in effect argues, is that extending that role to x_1 would trigger a chain reaction if matching is non-transitive. For all we know, x_1 is an element of a sorites series x_1, x_2, \dots, x_n where each x_i matches x_{i+1} , but where x_n does not match o . By parity of reasoning, we should extend the extension-fixing role right through the series to x_{n-1} , and then judge, in accordance with the First Rule, that x_n has the shade demonstrated in o , D_o . Yet the discriminability of x_n from o implies, again by the First Rule, that x_n does not have D_o . So we would have a contradiction. McDowell concludes:

Something which counts as a having a shade does not thereby count as a sample of that shade, bringing into the shade's extension anything indiscriminable from it... The status of a sample, a determinant of the extension of the concept in the relevant use, must be reserved for the original sample[.] (1994: 171)

¹¹ I am adopting the index notation from Brewer (1999: 175) to help make explicit what the original object for a given demonstrative concept is. Brewer's use of the notion is illustrated below in the text.

In a similar vein, Brewer writes.

Given a colour sample, A, let 'that_A shade' be an expression for the demonstrative colour concept grasped in virtue of a normal person's confrontation with A, the way of thinking of its colour which is provided by her experience of A. *Now suppose that something counts as having that_A shade if and only if it is indiscriminable in colour from A. ... it does not follow from any of this that a sample counts as having that_A shade if it is indiscriminable in colour from something, other than A, which counts as having that_A shade for the reason given. (1999: 175)*

The italicised rule here is clearly a notational variant of McDowell's First Rule.

Notice that if the First Rule were sound, it would provide a response to the sceptical challenge outlined at the end of the last section. That challenge granted that I am in a position to know that an object x matches an original sample o . It argued, however, that since matching is insufficient for identity of visible shade, I still am in no position to know that x has the specific shade o visually presents, since, for all I know, it doesn't. Yet, on the First Rule, the fact that x matches o in colour guarantees that x has D_o . Hence if I know that x matches o in colour, I would in fact be in a position to know that x has D_o too. Re-identification appears not to be beyond our epistemic powers, at least in the best synchronic cases. McDowell's and Brewer's treatment is thus at least *prima facie* promising.

2. The First Rule: a difficulty

Let's consider a minimal sorites series of just three objects a , b , c where a matches b in colour, b matches c , but a fails to match c . We can sum up the implications of the First Rule for this series in this table.¹²

	D_a	D_b	D_c
a	✓	✓	
b	✓	✓	✓
c		✓	✓

Notice two things here. First, b has both D_a and D_c , even though a is discriminably different in colour from c . Could this really be coherent given that these demonstrative concepts are to represent the most specific shades we are able to see things as having? Is it not like saying that a certain needle has both the length of a certain nail and the length of certain pin even though the nail is discernibly longer than the pin? The latter claim might appear to be incoherent. Second, notice that both a and c have in common D_b , even though they discriminably differ in colour. How can this be when the demonstrative phrase 'that_b shade', expressive of D_b , picks out the most specific shade we can see things to have?

Yet it may be too soon to dismiss this distribution of shades straight off as incoherent.¹³

Consider the following analogy. We have three monochrome objects, one red, one

¹² The table is inspired by Dokic & Pacherie (2001), whose article I came across after the first draft of this material was complete. They remark (as I will presently in the text) that the distribution of shades is *prima facie* odd, but do not argue (as I will) that there are objections of principle to it.

orange and one yellow. Similarly, we have three colour concepts *reddish*, *orangeish* and *yellowish*. We may assume that the objects fall under these concepts as shown.

	reddish	orangeish	yellowish
the red object	✓	✓	
the orange object	✓	✓	✓
the yellow object		✓	✓

This assignment is coherent because reddish, orangeish and yellowish are mutually overlapping shades of colour. (I will say a shade of colour is overlapping if two non-matching objects may have it.) Now, could not the demonstrative notions D_a , D_b and D_c be like *reddish*, *orangeish* and *yellowish* in representing mutually overlapping shades?

I will argue, however, that if the demonstrative concepts functioned as *reddish* etc. they could not fulfil the role the conceptualist needs them for. They would not allow him to account satisfactorily for the fine-grained structure of experience. The reason is that demonstrative concepts functioning in that manner of *reddish* etc. would not explain the possibility of the most fine-grained *changes* in colour that we can perceive.¹⁴

Suppose you are looking at the object a , and that while you are looking at it, it is instantaneously interchanged for c . We may indeed think of a and c as two stages of one and the same continuously perceived object o . Now, c by assumption does not match a in colour (we are presuming constant perceptual conditions). Thus when c is substituted

¹³ I am indebted to Tim Williamson for helping me to recognise this point.

¹⁴ I owe the suggestion that change perception may pose a problem for the First Rule to Ralph Wedgwood. I am unsure, however, whether he would approve of the way I develop the suggestion here.

for *a* you should visually experience the object *o* as changing in colour. Notice that you perceive the *object* as changing in colour. It is not just that there is a change in the *way* in which you perceive the object's colour. To get a grip on this distinction, consider the switch from seeing a four-sided figure as being a regular diamond to seeing it as being square.¹⁵ There is a change in the way you perceive the shape of the object here, all right. But you do not perceive the object as changing in shape. You do not, after the gestalt switch, see it as having a shape which it did *not* have before. By contrast, when you perceive the object as changing in colour, you *do* perceive it as having a colour which it did not have before.

These intuitive formulations point to an underlying principle of change perception. It is this:

Change-perception: If you perceive an object *o* as changing at *t*, then you represent *o* to have a certain property *F* after *t* which is incompatible with a certain property *G* which you represented *o* to have before *t*, where *F* and *G* are incompatible just in case nothing could be both *F* and *G* at the same time.

I will assume that this is a correct principle of change perception. It follows that we cannot explain your perceiving *o* to change in colour when *c* is substituted for *a*, if we suppose their colours are given merely as D_c and as D_a , respectively, *and* we suppose that these demonstrated colours are related as yellowish and reddish. For assume, by analogy, that you are fed reports by an observer who is looking at some strange animal outside your field of view. Until noon, he keeps on telling you 'it is reddish.' Then, after

¹⁵ See Peacocke (1992: ch. 3) for illustrations and discussions of this contrast.

noon, his report changes to ‘it is yellowish.’ Assuming these reports are perfectly reliable, could you tell, on the basis of them, whether the strange animal had changed in colour at noon? Clearly not.¹⁶ The animal might have been orange throughout. The only difference could have been that the observer subsumed its colour under *reddish* until noon, but for some reason took to subsuming it under *yellowish* thereafter.

The conceptualist might grant the conclusion that we cannot explain your perceiving *o* to change in colour, when *c* is substituted for *a*, if we suppose that their colours are given merely as D_c and as D_a , respectively, and that these demonstrated colours are related as yellowish and reddish. Yet he might resist the conclusion that the First Rule attributes too coarsely individuated shades. For he might take issue with our supposition that *c* is seen *merely* as having D_c and *a* seen *merely* as having D_a . The conceptualist takes his cue here from our intuitive formulation above, when we said that ‘when you perceive the object as changing in colour, you perceive it as having a colour which it did *not* have before.’ In light of this he takes himself to have two options. One is to say that *a* before being substituted for *c* is seen not only as D_a but also as not- D_c . The other is to say that *c*, after the switch, is seen not only as D_c but also as not- D_a . Of course, nothing can have both D_c and both D_a and not D_c , just as nothing can have both D_a and both D_c and not D_a .

The first option is not tenable since the concept D_c is not available until *c* comes into view. What, however, can we say about the second option? To rule that one can never

¹⁶ I am assuming that the grounds for your judgement are restricted to the semantic content of the reports, and so do not include any conversational implicatures of the reporter shifting from formulating his reports one way to his formulating them another way. Thanks to Brad Skow for alerting me to the need for this restriction.

see an object as not having a certain shade of colour may be too strong. Perhaps we can sometimes do so. Yet no plausible model of how or why we see c as not D_a seems to be available to the conceptualist. One candidate model of how we can see an object as not having a certain shade of colour is that we see it as having a positive colour which excludes its having the colour in question. This model is unavailable to our conceptualist since the idea that we see c as not D_a was introduced precisely in order not to have to appeal to a more specific positive colour, incompatible with D_a , which we see c as having. Another candidate model of how we can see an object as lacking a certain shade of colour is this. Arguably, there is such a thing as attentively ‘hunting’ for a certain nuance in a seen surface or expanse. It may be that, if I launch myself into such an attentive hunt, and the hunt yields no positive result, this failure manifests itself phenomenologically in my seeing the pen as lacking the hunted-for shade.¹⁷ Again, this model does not give our conceptualist a plausible account. The proposed mechanism involves an attentive hunt, which is an intentional mental activity. The idea that, in order to perceive the posited continuing object o as changing in colour here, one *needs* to engage in such activity, lacks plausibility. Indeed, what would prompt our subject to such activity, supposing she did engage in it? A plausible answer would be that she experienced something change in colour. But a plausible answer of this sort is unavailable to the conceptualist who appeals to the second model precisely to explain change perception in the first place.

Perhaps the conceptualist will propose other ways of explaining the relevant change perception consistently with the First Rule. I for my part cannot see any reason to think these other proposals will fare better than the ones we have considered. I conclude, then,

¹⁷ John Campbell suggested this second model to me (in conversation).

that if the conceptualist is to vindicate his commitment to demonstrative concepts being ‘exactly as fine-grained as the experience’, as McDowell says they are, he is bound to give up the First Rule in favour of a rule setting stricter application conditions for the demonstrative concepts. I now turn to consider such a stricter rule and its consequences.

3. The Goodmanian Rule and the sorites sceptical argument

Consider an arbitrary object x . What does it have to be like to be subsumed under D_o if we are to avoid the problems discussed in the last section? Clearly x must match o in colour. But this is not sufficient. For the fact that x matches o is consistent with there being another object y which matches o but not y . That tells us that D_o cannot apply to both x and y ; this was the lesson of the last section. As o *ex hypothesi* matches both x and y it seems to be gratuitous to rule that D_o applies to one of them but not the other. So it seems plausible to conclude that D_o applies to neither. If so, it follows that D_o applies to an object x only if x matches just the same objects as o matches. This relational condition of matching just the same objects is a transitive one; indeed, it is an equivalence relation. So this necessary condition is surely also sufficient. We have here arrived at what I will call the Goodmanian rule for demonstrative concepts, as it is structurally parallel to Goodman’s identity criterion for qualia.

The Goodmanian Rule: An object x has D_o iff for all objects z , z is indiscriminable in colour from x just in case z is indiscriminable in colour from o .

By adopting the Goodmanian Rule, it appears that the conceptualist can account for the fineness of grain of experience and meet the requirements of the principle of change perception, without coming into conflict with the non-transitivity assumption. However,

he can do so only at the cost of undermining the recognisability requirement. Once we give up the First Rule in favour of the Goodmanian Rule, the sceptical worry about the possibility of knowledgeable re-identification is back in force. The intuitive worry is this. On the Goodmanian rule, indiscriminability of shade between x and a target object o is insufficient for x to have the specific shade demonstrated in o . So even assuming (as we do) that you may know that x matches o in colour, how could you ever know that x has the colour o visually presents, D_o ? For all you know, something matches x which does not match o , or vice versa. Hence, for all you know, x is not D_o .

One line of reply to the sceptical worry is this. When you perceive x you are in a position to tell (we are assuming) that x matches o in colour. Granted, knowing so much does not put you in a position to rule out the possibility of there being an object which matches just one and not the other. Still, your perception of x may provide *defeasible* or *prima facie* justification for believing that x has D_o . Of course, this justification would be defeated if you came across an object which matches one but not the other of the two. But in the absence of such an encounter, your *prima facie* justification is undefeated, and undefeated defeasible justification is justification *tout court*. Suppose moreover x in fact matches just the same objects as o . Then the judgement that x has D_o is not only justified but true. There is no reason to say that it must be only accidentally true in the way characteristic of Gettier cases. So there is no reason to deny that this judgment cannot, in favourable cases, amount to knowledge. Or so the current anti-sceptical reply says.

However, the possibility of sorites series presents an obstacle to this anti-sceptical argument. The anti-sceptical argument assumes that any defeasible, knowledge-

underwriting, justification for the (we-assume-true) belief that x is D_o provided by the experience of x , where x visually matches o in colour, remains in force in the absence of a positive encounter with a non-common third party. But that assumption is questionable. It may well be that x and o are the first pair in a sorites series o, x_1, \dots, x_n where each object matches the next yet where the last, x_n , fails to match o . By the argument from defeasible justification, you would be defeasibly justified in believing each of the ascriptions

(K₁) x_1 has D_o

(K₂) x_2 has D_{x_1}

(K₃) x_3 has D_{x_2}

...

(K_n) x_n has $D_{x_{(n-1)}}$

Now, since x_n does not match o in colour, and since, as we are assuming, you are in position to know this, you are in a position to know that

(N) x_n does not have D_o .

Moreover, we may assume that you have two pieces of conceptual and logical knowledge. First, supposing that your usage of demonstrative concepts is governed by the Goodmanian rule, we may assume that you know, at least implicitly, that ascriptions of demonstrated shades are subject to the principle-scheme

(G) x has $D_o \leftrightarrow \forall z (\text{Matches}(z, x) \leftrightarrow \text{Matches}(z, o))$

Second, if you understand the universal bi-conditional, you may assume that you know this relation is transitive. That is, we may assume that you know this principle-scheme, at least implicitly

(B) $[(\forall x(F_1x \leftrightarrow F_2x)) \ \& \ \dots \ \& \ (\forall x(F_{n-1}x \leftrightarrow F_nx))] \rightarrow \forall x(F_1x \leftrightarrow F_nx)$

The contrapositive of (B) is:

$$(B^{cp}) \neg \forall x (F_1 x \leftrightarrow F_n x) \rightarrow \neg [(\forall x (F_1 x \leftrightarrow F_2 x)) \& \dots \& (\forall x (F_{n-1} x \leftrightarrow F_n x))]$$

An instance of (B^{cp}) for (G) is

$$(B^G) \neg \forall z (\text{Matches}(z, o) \leftrightarrow \text{Matches}(z, x_n)) \rightarrow \neg [(\forall z (\text{Matches}(z, o) \leftrightarrow \text{Matches}(z, x_1)) \& \dots \& (\forall z (\text{Matches}(z, x_{n-1}) \leftrightarrow \text{Matches}(z, x_n)))]$$

Substituting in (B^G) on the basis of (G) we get

$$(S) x_n \text{ does not have } D_o \rightarrow \neg (x_1 \text{ has } D_o \& x_2 \text{ has } D_{x_1} \& \dots \& x_n \text{ has } D_{x_{(n-1)}})$$

We may assume then, that you know (S) , at least implicitly.

At this point we need to invoke an assumption of closure of empirical knowledge, and of warrant, under known entailment. Closure assumptions are of course not uncontroversial; however, the default view must be that knowledge and warrant are transmitted by known entailment.¹⁸ Moreover, conceptualists like McDowell and Brewer seem to be happy to go along with that assumption.¹⁹ Now, assuming that you are warranted in believing each claim (K_i) , and you master the rule of $\&$ -introduction, you should be warranted in believing their conjunction

$$(K) x_1 \text{ has } D_o \& x_2 \text{ has } D_{x_1} \& \dots \& x_n \text{ has } D_{x_{(n-1)}}$$

However, as you know (N) and (S) , you may knowledgeably deduce

$$(\neg K) \neg (x_1 \text{ has } D_o \& x_2 \text{ has } D_{x_1} \& \dots \& x_n \text{ has } D_{x_{(n-1)}})$$

Your knowledge that $(\neg K)$ clearly destroys your warrant for believing (K) . Ipso facto, it makes it the case that you are no longer defeasibly justified in believing each claim (K_i) . Could one argue that this situation could be remedied by giving up on some of these claims (K_i) , and in consequence become defeasibly justified in believing a select subset

¹⁸ For rejections of closure, see Dretske (1970) and Nozick (1981).

¹⁹ Brewer (1999: ch.2) explicitly distances himself from anti-sceptical views that deny closure. In a passing reference, McDowell (1982) is officially agnostic about the closure assumption. It is notable, however, that he does not seek to address sceptical worries by rejecting it.

of the claims (K_i)? The problem with this is that your prima facie warrant for accepting each (K_i) claim may well be equally good under the circumstances. If, but only if, there had been one such claim for which your warrant was worse than that for all the others, it seems you could rationally respond to your knowledge that ($\neg K$) by rejecting that claim in particular. You could rationally single out that claim as the culprit responsible for the falsity of the conjunction. However, there is no reason, in principle, to think that there need be such a ‘worst off’ claim. The upshot, I suggest, is that the presumed availability of sorites series acts as a standing defeater of any transition from known matching in colour to purported recognition of demonstrated shades.²⁰

²⁰ This sorites-based argument can be simplified if we make the uniqueness assumption that a perceivable object has one and only one Goodmanian shade. On this assumption, we can formulate the Goodmanian Rule as stating

the shade of $x = D_o \leftrightarrow \forall z (\text{Matches} (z, x) \leftrightarrow \text{Matches} (z, o))$

Since indiscriminability is reflexive, we have that

the shade of $x = D_x$

Hence, by Leibniz Law

$D_x = D_o \leftrightarrow \forall z (\text{Matches} (z, x) \leftrightarrow \text{Matches} (z, o))$

If we assume that our use of demonstrative concepts is governed by this criterion, the sorites-based argument takes the following form. Consider a colour sorites series x_1, \dots, x_n subject to the usual conditions: each x_i is indiscriminable in colour from x_{i+1} but x_1 is discriminable in colour from x_n . Then we have that you are equally justified in believing each identity

$D_1 = D_2$

$D_2 = D_3$

...

$D_{n-1} = D_n$

Supposing that you know the transitivity of identity and contrapositive reasoning, you know as a logical truth that

$D_1 \neq D_n \rightarrow \neg (D_1 = D_2 \ \& \ \dots \ \& \ D_{n-1} = D_n)$

You know empirically that

$D_1 \neq D_n$.

Hence, assuming closure of empirical knowledge over known logical entailment, you know empirically that

$\neg (D_1 = D_2 \ \& \ \dots \ \& \ D_{n-1} = D_n)$

From this point, the argument proceeds as in the text.

It will not have escaped the reader's attention that the structure of the above is parallel in one respect to the lottery paradox. The anti-sceptical reply we have been examining is analogous to the view that one might have a sort of defeasible justification for believing that a given ticket *o*, out of a certain number *N* in a fair lottery, is a non-winning ticket, sufficient for that belief (if true) to constitute knowledge. Now, since there is need be nothing special about the given ticket *o*, this anti-sceptical argument can be repeated for each ticket in the lottery. So you know of each ticket except for the one that wins that it is a non-winning ticket. This has struck a number of philosophers as absurd, and I believe rightly so. To return this verdict is not to deny, of course, that it is delicate matter to spell out the reasons for the verdict in a way that allows for satisfactory overall epistemology. (If this had been easy, the lottery paradox wouldn't have been a paradox.) If this is our preferred verdict about the lottery case, as I believe it ought to be, the conclusion should be the same about the anti-sceptical suggestion about the whether we can know that *x* has *D*₀. Furthermore, the anti-sceptical suggestion in the phenomenal sorites case is in another respect not parallel to the anti-sceptical strategy in the lottery case, but considerably shakier. For whereas the lottery case may be so specified that non-winning tickets are (to put it mildly) much more frequent than winning tickets, there is no reason at all to believe that matching pairs meeting the Goodmanian condition are more frequent than matching pairs not meeting it – indeed, the opposite would seem plausible. I conclude that recognition of shades becomes impossible on the Goodmanian Rule.

This sceptical argument is clearly not limited to just diachronic cases of purported recognition. Hence it undermines even the weak, synchronic, merely cross-modal recognition condition. I do not see any room for an even weaker constraint on what it

takes to have a recognitional capacity. The conceptualist, then, seems to be committed choose between these options: (i) deny non-transitivity; (ii) deny that demonstrative concepts are subject to a recognition requirement; (iii) deny that the concepts of fine-grained features we need to possess to experience something as having them are demonstrative concepts. The first option is *prima facie* implausible, and is likely merely to move the problem rather than do anything to solve it; unfortunately, going into this option in detail would take us too far afield at this point, and must be left for another occasion.²¹ Setting it aside, we are in a position to draw the following conclusion: if the concepts of fine-grained features needed to experience something as having them are demonstrative concepts, then demonstrative concepts are not subject to a recognitional requirement.

²¹ I will however offer the following brief comment. If we affirm transitivity, the underlying problem addressed in this paper is likely to move in the following way. In this present paper, we have been assuming, for the sake of argument, that one is in a position to know that *x* and *y* match in colour. We have also been assuming that there is no challenge in principle to one's capacity to know that the relevant parameters of the conditions of perception are the same when one compares *x* and *y* for colour-match, and when one compares *y* and *z* for colour match. Granting these two 'epistemically optimistic' assumptions, we have raised problems for the claim that one is in a position to know that *x* and *y* have the same shade. Now, the standard strategies that are deployed in arguing for the transitivity of matching have the consequence of casting serious doubt on one or both of these optimistic assumptions. For example, Raffman (2000) and Graff (2001) argue that the impression that matching is not transitive arises from a failure to take into account certain changes in the conditions of perception when one shifts from comparing *x* and *y* for match to comparing *y* and *z* for match. If *x* and *z* do not match when *x* and *z* are compared, then the shift from comparing *x* and *y* to comparing *y* and *z* *inevitably* involves certain relevant changes in the parameters of the conditions of perception. Or so Raffman and Graff suggest. If this is right, our assumption that one is in a position to know that the relevant parameters of the conditions of perception are the same of course becomes all the more problematic.

ESSAY 3

Conceptualism, Self-Conscious Thought, and Systematicity

0. Introduction

Philosophers have long been interested in the connection between perceptual experience and thought. One way of addressing one aspect of that connection is by exploring the extent to which perceptual experience itself is a concept-demanding affair (as thought presumably is). Interest in this latter question is of course nothing new either, as a perusal of the views of Kant and Kuhn, to name but two, quickly makes clear.¹ However, over the last couple of decades, it has spawned a particularly lively debate between so-called conceptualists and nonconceptualists. In this paper, I critically examine conceptualism, and argue that conceptualists are committed to give up two principles about concepts that have been widely assumed, not least by conceptualists themselves. I will call these the Kantian Principle and the Systematicity Principle. The first says, roughly, that concepts essentially are exercisable in self-conscious thought; the second, roughly, that they essentially are exercisable in cognitive states exhibiting a certain systematicity or re-combinability. I will not argue that this refutes conceptualism, since I will not argue that either of the principles is non-negotiable. I

¹ Kant famously remarked that 'intuitions without concepts are blind' (1787: A51/B75). On some interpretations of Kant, e.g. that in McDowell (1998a), this remark epitomises a view on which the use of concepts is inextricably implicated in sensory experience of the world. (For a contrasting interpretation on this point, see Hanna (2005).) Kuhn (1962) argued that a conceptually structured paradigm heavily influences what we see when we look at things. Presumably, this is one of the features of paradigms that underlies his assertion that people with different paradigms live in different worlds. The fact that both Kant and Kuhn can be cited as examples of views on which the use of concepts is tightly interwoven with sensory experience of the world goes to show how deeply this broadly conceptualist view cuts across the question of sense perception's objectivity versus its possible cultural relativity.

conclude, however, that the conceptualist is left with two options: either to take the radical tack of arguing that not even thought exhibits systematicity, or face an embarrassing question over what the principled distinction between conceptual and nonconceptual representational capacities for her consists in.

Conceptualism, for the purposes of this paper, I take to be the following thesis:

Conceptualism To have a perceptual experience as of ...F..., a subject needs to possess a concept of F/F-ness/F-hood/...

The slashes and dots figure at the end here to allow that different instances of this schema have to be formulated slightly differently, depending on the semantic type of the expression that takes the place of 'F'.² I will disregard this complication in what follows. To note one instance of the schema: in order to have a perceptual experience as of something being red, a subject needs to possess a concept of red.

Conceptualism, in this sense, is usefully seen as a generalisation to perceptual experience of a constraint that seems to hold across a wide range of intentional states.

Most notably, for belief:

Conceptualisation Constraint – Belief To have a belief that ...F..., a subject needs to possess a concept of F.

² See Introduction, section 3.1, for more on this.

The corresponding constraint for desires, hopes, and other paradigmatic propositional attitudes seem no less plausible than this constraint for belief. Another plausible instance is admiration:

Conceptualisation Constraint – Admiration In order to admire F, a subject needs to possess a concept of F.

For example, in order to admire ballet dancers, one needs to possess a concept of ballet dancers.³ The plausibility of this last constraint is interesting, in that it shows that the plausibility of the conceptualisation constraint is not confined to propositional attitudes. For, as Tim Crane (2001: 112-4) and Graeme Forbes (2006: 52-68) have recently argued, admiring ballet dancers is not a case of a propositional attitude, but rather of what we may call an objectual attitude. In generalising the conceptualisation constraint to perception one need not tacitly assume, as is sometimes suggested, that perceptual experience is a propositional attitude. Perhaps one needs to assume that perceptual experience is an intentional state of mind; yet, in a generous sense of intentional, I take that to be highly plausible.

³ One might object as follows. ‘Consider a desert dweller who has no notion of ballet, ballet dancing, or ballet dancers. One day he encounters a stranger in the desert who does something magnificent (wholly unrelated to ballet) and whom he consequently comes to admire. As it happens, the stranger is a ballet dancer. So our desert dweller admires a ballet dancer.’ What this objection shows is that conceptualisation constraint for admiration is true only if ‘admires’ is taken on the intensional, and not on the extensional (also known as the relational), reading. On the intensional reading, one can admire ballet dancers even if there are none, and admire a ballet dancer without admiring any ballet dancer in particular. Moreover one can admire ballet dancers without admiring people with a certain arcane genetic pattern, even if all and only ballet dancers have the relevant arcane genetic pattern. On this reading, it does not follow from the fact that our desert dweller admires a magnificent stranger that he admires a ballet dancer, even if the stranger who provoked the former admiration was, in fact, a ballet dancer.

Conceptualism in our sense has been defended in the last two or three years by Alva Noë (2004), Jeff Speaks (2005) and anew by John McDowell (forthcoming).⁴ This is one reason critically to examine the view again. Another is that doing so brings to the fore interesting questions about the connections between the possession of concepts and self-conscious thought, on the one hand, and between the same and systematicity, on the other, as anticipated. Thus, more fundamentally, to examine conceptualism is a way of bringing into relief the underlying question of what the principled distinction between conceptual and nonconceptual representational capacities may be, if indeed there is any such.

A final word of warning before I proceed. ‘Conceptualism’ is often taken as a label for the thesis that perceptual experiences only have conceptual content, or that they have the same kind of content as thoughts. Either thesis may well be related in interesting ways to conceptualism in our sense, but should not be confused with the latter.⁵

The structure of the paper is this. In the next section, section 1, I rehearse the objection to conceptualism from the specificity, or fineness of grain, of experience. I also criticise two standard replies to the objection, adverting respectively to recognitional and demonstrative concepts. In section 2, I consider two arguments for the claim that lexical, recognitional, and demonstrative concepts exhaust the candidates for being those possession of which is needed to experience something as having a certain specific feature – being red₁₂₃, say. Both of these arguments assume the Kantian

⁴ Earlier notable defences of conceptualism include Armstrong (1968), McDowell (1994), and Brewer (1999).

⁵ Conceptualism in our sense is roughly equivalent to what Byrne (2005) calls ‘state conceptualism’ and Speaks (2005) calls ‘relative conceptualism’. See the Introduction, especially sections 6 and 7, for remarks upon the relation between conceptualism in this sense and the ‘content conceptualist’ claim that perceptual experiences have the same kind of content as thoughts.

Principle. In light of our conclusions in section 1, I argue that they show the conceptualist to be committed to give up the Kantian Principle. In section 3, I argue that an appeal to systematicity cannot do what the conceptualist may want it to do: viz. to serve (in part) to mark off conceptual from nonconceptual representational capacities in a principled way. In section 4, I respond to some objections.

1. The objection from fineness of grain, and some responses

A natural place to begin the discussion of conceptualism is with a classic purported counterexample to it, noted in the late seventies by David Kelley (1980) and Gareth Evans (1982: 229) and pressed by many later critics of the view. The would-be counterexample trades on the specificity or fine-ness of grain of perceptual experience. Consider a subject, Sam, who looks at a piece of paper having a highly specific shade of red – let's call it 'red₄₅' – under good conditions in which that shade is clearly visible to him. So Sam has a visual experience as of something being the colour red₄₅. But he does not possess a concept of red₄₅, or, at least, *need* not possess such a concept.

Now some of the concepts a thinker possesses at a given time are such that she is a competent user of a word expressing that concept. Call such concepts 'lexical concepts'. Among the lexical concepts, some are such that the thinker is a competent user of a *non-demonstrative* word expressing that concept. Call such concepts 'non-demonstrative-lexical concepts'. If *all* concepts possessed by thinkers were non-demonstrative-lexical concepts, the would-be counterexample would clearly stand. People do not and certainly need not be competent users of non-demonstrative words for all the shades of colours they see things as having. However, as conceptualists quickly pointed out, not all concepts are non-demonstrative-lexical concepts. Some are so-called recognitional concepts: concepts we possess at least in part thanks to a capacity to recognise things

falling under the concept when we are presented with them. Perhaps Sam in our example is bound to possess a recognitional concept of red₄₅, even if it is a concept that relies on a recognitional capacity which lasts only for a short time. John McDowell (1994: 57-60) made this suggestion in his first response to Evans.

It evoked strong criticism from Diana Raffman (1995), Christopher Peacocke (2001), and Sean Kelly (2001) among others. These critics argue that, as an empirical matter, our recognitional capacities across time cut more coarsely than the shades of colour we can synchronously differentiate in experience at a time. This reflects an inbuilt limitation on our memory capacities. Therefore, even if these recognitional capacities are supposed to last for only a very short while (as McDowell underscores), they will be more coarse-grained than the experience. If red₄₅ is about as specific a shade of colour as Sam can synchronously differentiate in experience, he will not have a recognitional concept of red₄₅ as opposed to a more generic or determinable shade of which red₄₅ is a determinate (along with, perhaps, red₄₄ and red₄₆ among others).

Even worse: even if the relevant recognitional capacity does not need to outlast the occurrent experience of an object as red₄₅, the subject does not have it. This is because the non-transitive nature of phenomenal indiscriminability underwrites a robust and convincing sceptical challenge to our capacity knowledgeably to judge, on the basis of experience, of another object looking red₄₅ to us, that it is red₄₅.⁶

Another type of concepts invoked by conceptualists are so-called demonstrative concepts, or more specifically perceptual-demonstrative concepts (for brevity, I use the

⁶ I develop an argument to this effect in essay 2.

term demonstrative concepts hereinafter). These are concepts we use when we think of object or properties we currently are perceiving, and think of them in a way that would most naturally be expressed in natural language in terms of demonstratives. If I see a number of people before me, I may think of one of them in a way that I would most naturally express by saying ‘*That* guy is so-and-so’. We may think of properties in an analogous way. Looking at an oddly shaped button on someone’s jacket, I may think to myself something I would naturally express by saying ‘I should get myself a button shaped *that* way.’ These concepts seem congenial to the conceptualist, allowing her to suggest that, when Sam has an experience of something as red₄₅, he is bound to possess a demonstrative concept of that shade, in terms of which he can think, for example, ‘I want a shirt coloured *that* way.’

Some objections to this demonstrative suggestion, in particular that pressed by Shaun Kelly (2001), rely on the assumption that demonstrative concepts are, in effect, a special case of recognitional concepts. This latter assumption is moreover one that some influential conceptualists have shared.⁷ However, the assumption that demonstrative concepts presume a recognitional capacity for the things they are concepts of is doubtful, as has been argued by Jeff Speaks (2005) and others. So objections relying on this assumption do not have much force against the claim that our subject, Sam, is bound to possess a demonstrative concept of red₄₅ when he experiences something as having that shade.

A better objection to that demonstrative claim relies on a link between demonstrative concepts and attention. This objection, which we may call the Objection from Attention,

⁷ See McDowell (1994: 56-60).

relies on two premises. The first is that grasp of a demonstrative concept presumes attending to the object or property the concept is a concept of. The second is that that perceptual experience as of a certain object having a certain property does not require attending to that object and to that property. If this is so, Sam does not after all *need to* possess a demonstrative concept of red₄₅ in his experiential situation, even if he of course *may* very well do so.

The two premises of this objection are both highly plausible, at least as far as reflection on our thought and experience is to be trusted on these matters. John Campbell (2002) has done much to make vivid the essential link between grasp of demonstrative concepts and conscious attention, affirmed by the first premise. Consider the case above where I see a number of people before me. How is it that I am able to think about one of them in particular when I think the thought expressible as ‘That guy is so-and-so’? Suppose my attention is disengaged, so that it is as if I am having a ‘sea of faces’ before me. If I succeed in thinking of any particular person in this condition, it certainly would not be by means of vision, and, consequently, not by a perceptual-demonstrative concept. If I am to succeed in understanding my own words as expressing a visual concept of someone, I must engage my attention, and, in Campbell’s phrase, ‘highlight’ one of the men in my experience. A parallel point about grasping demonstrative concepts of properties seems no less plausible. Suppose my interlocutor says, gesturing at someone, ‘One wouldn’t want to be that way’. How am I to understand what way my interlocutor says here? Is he making a remark about the person’s posture, his bodily shape, his apparent mood, or something else? And what determines which? It seems highly plausible that at least part of what determines this is what my interlocutor attends

to. I won't grasp a demonstrative concept like the one he has unless I visually attend to the same apparent aspect of the man before us.

The second premise of the objection from attention is that it is not essential to consciously attend to an object or property in order to have a perceptual experience as of that object being that way. As far as introspective reflection is to be trusted, this seems overwhelmingly plausible. Suppose that I stretch both of my arms out horizontally in front of me, palms up and towards me. When I am now attending to my right hand, it seems undeniable that I still have visual experience of my left hand as being a certain way. Certainly, my visual experience of the left hand is not in some sense as 'informative' as that of my right hand, but I still *have* visual experience of it. If I were to shift attention from the right hand, it is well-nigh incredible that I would ipso facto cease to have experience of it. When we turn to selective attention to properties, the same point seems if anything even more forceful. Going by introspection, it seems I can selectively attend to the movement of my hand, or to the pattern of skin-tones over it, and that I can do the one without the other. Again, it seems incredible that, were I selectively to attend to the movement and not to the colours, my hand would cease to look any colour to me – I would cease to have an experience of my hand as having a certain colour.

How have conceptualists responded to this objection? Not, as far as I know, by denying the link between demonstrative concepts and attention. As far as the problem has been considered, the chosen target has been rather the denial, in the second premise, that perceptual experience of something presumes attending to it. Alva Noë (2004) is probably the clearest case of a conceptualist who takes this line. His argument relies

heavily on work done by experimental psychologists on the phenomena of change blindness and inattentional blindness.⁸ He takes this work to support the drastic conclusion that we really don't actually enjoy perceptual experience of the things we don't attend to. The reason we think otherwise is *not*, on Noë's account, that introspection as such is misleading. The reason is that we – and here 'we' means philosophers of perception first and foremost – stubbornly put a fallacious spin on what introspection purports to let us know. What introspection purports to let me know, and indeed does let me know, when I now stare at my right hand, is that I very, very easily *could* have had an experience of my left hand, viz. by turning my attention to it. The fallacious spin is that I *actually* have experience of my left hand. In reflecting on their own experiences, philosophers of perception have in other words been massively prone to the fallacy of inferring actuality from easy possibility.

The assessment of Noë's view would lead us into many controversial empirical and philosophical issues, and rapidly take us too far afield. I will therefore limit myself to recording that I do not find the evidence cited by Noë strong enough to warrant the denial of the second premise of the Argument from Attention, particularly not in the presence of other, putatively conflicting evidence, such as George Sperling's (1960) experiments on sensory memory.⁹ In particular, even if the evidence from change blindness and inattentional blindness supported the idea that we do not have experience of *objects* we do not attend to (and I am sceptical even of that), it is a much stronger, and much less plausible claim which has it supporting the idea that we lack experience of *properties* we do not attend to. Surely, we experience a stop-sign as having a certain

⁸ See Rensink, O'Regan & Clark (1997) on change blindness and Mack & Rock (1998) on inattentional blindness.

⁹ Tye (2006a) draws attention to the relevance of Sperling's experiments to the issue of whether we have experience of things beyond what we attend to.

shade of red (it looks some shade of red to us) even when we attend selectively to its shape.¹⁰

I conclude then that the conceptualist's appeal to demonstrative concepts, like her appeal to recognitional concepts, does not defuse the objection from fineness of grain.

2. Concepts and self-conscious thought

The conclusions of the last section might seem to leave the conceptualist in a bad place, lacking a reply to the purported counterexample from fineness of grain. However, on reflection, and looking at things with more initial sympathy for the conceptualist, things are perhaps not so bad for her after all. Even if the conceptualist has not refuted the counterexample, what has her opponent done to establish it as genuine? What her opponent really has done is to take a couple of steps towards an argument by exclusion. Let's call the concepts of F one needs to possess to have an experience of something as F, if conceptualism is right, 'perceptual concepts'. What has been argued is that perceptual concepts are not non-demonstrative-lexical, recognitional or demonstrative concepts. But why should we believe these to exhaust the candidates for being perceptual concepts? Interestingly, the list has been taken as exhaustive not only by the critics of conceptualism, but also conceptualists themselves. At any rate, other candidate concepts for this role have not, to my knowledge, been mooted. Yet there is little *explicit* discussion or defence of this shared presumption to be found in the literature.

In this section, I will set out two closely related lines of defence of the presumption that the three candidates exhaust the candidates for being the perceptual concepts. I in no

¹⁰ Lerman (ms.) provides further discussion and defence of this claim.

way intend to suggest that these lines of argument are original; I take them, rather, to make explicit motivations that are more-or-less implicit in several writers on these issues.

The first line of argument we may call the Argument from Expression. I will state it first in a somewhat bold version; some possible refinements will be noted below.

The Argument from Expression

- (1) Assume that conceptualism is true, and so that relative to a experience as of ...F..G..., there are certain concepts of F, of G, one needs to possess to have that experience. Call these 'perceptual concepts'.
- (2) In order for someone genuinely to possess a concept, he must be able to exercise that concept in thinking self-conscious thoughts, that is, in thinking thoughts that are available for immediate self-ascription (whether or not they are actually so self-ascribed). (**The Kantian Principle**).
- (3) So perceptual concepts must be exercisable in self-conscious thoughts. (From (1) and (2).)
- (4) Any self-conscious thought that can be entertained by a rational subject is such that, if the subject were linguistically competent, she could entertain the very same thought, using the concepts she actually does, and could linguistically express it, although the linguistic expression in question may be ineliminably demonstrative. (**The Expression Principle**)
- (5) So perceptual concepts must be exercisable in thoughts such that, if the subject were linguistically competent, she could entertain the very same thoughts, using the perceptual concepts she actually does, and could linguistically express them,

although the linguistic expression in question may be ineliminably demonstrative. (From (3) and (4).)

- (6) Now, take an arbitrary perceptual concept, and consider the linguistic expression of one of the self-conscious thoughts that are the best candidates for being those in which the perceptual concept is used. There is a (semantically unitary) part of this linguistic expression that expresses the contribution made by the perceptual concept to the self-conscious thought in question; for short, there is a part of it that expresses the perceptual concept. (Compositionality)
- (7) Either the relevant part of the linguistic expression is a non-demonstrative expression, in which case the perceptual concept is a non-demonstrative-lexical concept – a concept such that the subject possessing it is a competent user of a non-demonstrative word expressing. Alternatively, the relevant part is a demonstrative expression, in which case the perceptual concept is a demonstrative concept (or perhaps a recognitional concept, if the demonstrative concept in question is subject to a recognition condition).
- (8) Therefore, perceptual concepts are either lexical or demonstrative (or recognitional) concepts.

The crucial steps of the Argument from Expression are steps (2) and (4). I will comment on them in turn.

I call the step (2) ‘the Kantian Principle’ as it is *broadly* reminiscent of Kant’s idea according to which the ‘I think’ of apperception must be capable of accompanying all my representations. It is clearly a consequential claim. Cognitive ethologists generally think that non-human animals, with at most a few exceptions such as chimpanzees, are

not capable of self-conscious thought. The Kantian Principle entails, then, that almost no non-human animals have concepts. If so, the conceptualisation constraint on beliefs and desires further entails that almost no non-human animals have beliefs or desires. This result evidently clashes with our firm common-sense conviction that a dog may think that a cat has run up a tree, or want to catch a bone thrown before it. The Kantian Principle also clashes with a widespread view in cognitive science, given a pretty plausible assumption. The widespread view is that subjects have tacit knowledge of things about which they are not capable of having self-conscious thoughts. For example, Chomskyan linguistics posits that eight-year old speakers tacitly know that such-and-such parts of sentences play the role of grammatical subject (say). The pretty plausible assumption is that tacit knowledge is knowledge. Given this, and the well-nigh uncontroversial idea that knowledge entails belief, it follows that eight-year-old speakers have the belief that such-and-such parts of sentences play the role of grammatical subject. The conceptualisation constraint for belief entails then that eight-year-old speakers possess the concept of grammatical subject.¹¹ Yet eight-year-old speakers are generally not capable of having reflective thoughts (thoughts available for self-ascription) about grammatical subjects. This situation contradicts the Kantian Principle.

Now, believers in the Kantian Principle will generally seek to soften these contradictions with common sense and cognitive science by allowing for *nonconceptual* representational capacities to which the principle does not apply. By virtue of such capacities, a dog may somehow quasi-believe that so-and-so and quasi-want that such-and-such. By virtue of such capacities, an eight-year-old may similarly quasi-know that

¹¹ Strictly, the claim that knowledge entails belief is not necessary to reach this conclusion, since the conceptualisation constraint is no less plausible for knowledge that P than it is for belief that P.

such-and-such parts of sentences play the role of grammatical subject; tacit knowledge is precisely such quasi-knowledge (*pace* the pretty plausible idea that tacit knowledge is knowledge). The Kantian Principle will thus typically be given a role of specifying part of what marks off, in a principled way, conceptual from nonconceptual representational capacities.

Interestingly, enthusiastic defenders of the Kantian Principle include such conceptualists as McDowell. (This is interesting in as much as conceptualists might have been expected to play down, rather than up, what it takes to have a concept.) He writes:

We would not be able to suppose that the capacities that are in play in experience are conceptual if they were manifested only in experience, only in operations of receptivity. They would not be recognisable as conceptual capacities at all unless they could also be exercised in active thinking, that is, in ways that do provide a good fit for the idea of spontaneity. (1994: 11)

McDowell takes it to be essential to what he calls ‘active’ thinking here that that is potentially self-critical thinking – thinking that ‘takes place under a standing obligation to reflect about the credentials of the putatively rational linkages that govern it.’ (1994: 12; see also 1994: 66-7) Such critical reflection upon one’s thoughts clearly presupposes that one has reflective access to what one is thinking. Thus McDowell is clearly committed to the Kantian Principle. We shall return below to the question whether conceptualists are wise to be so committed.

The claim in step (4), the Expression Principle, embodies the idea that self-conscious thoughts are somehow essentially linguistically articulable, although perhaps articulable

only in essentially demonstrative (or otherwise indexical) terms. It is important to stress the qualification about the possibly essentially demonstrative character of these expressions, lest the principle will seem 'over-intellectualised'. At the limit, the linguistic expression of a self-conscious thought may take simply the form '*That* thing is *thus*.'¹² So the principle does not entail that thinkers of self-conscious thoughts must be capable of giving them a highly articulated expression, in terms understandable to someone not having had any similar experiences.

The Expression Principle, as here formulated, does not entail that any self-conscious thought on the part of any subject is such that the subject can linguistically express it. The Principle is thereby consistent with the claim that aphasiacs or (perhaps more generally) sufferers from localised language production impairments can have self-conscious thoughts. It is also consistent with the perhaps more contentious claim that pre-linguistic infants can have self-conscious thoughts. At any rate, it is consistent with these claims if aphasiacs or pre-linguistic children could entertain the very same thoughts were they to become linguistically competent. Note also that the linguistic expressions need not be overt. They need not be cases of talking or writing, but simply of saying something to oneself.

Some conceptualists, notably McDowell (1994: 165), have been happy to accept the Expression Principle. When it comes to whether anyone ought to do so, it is useful to note that there are ways for the principle to fail that do relatively little damage to the spirit if not the letter of the Argument from Expression. Suppose the principle fails

¹² According to Brewer (1999: ch. 6), utterances of the form 'That is thus' are indeed the most straightforward or canonical ways of expressing the thoughts whose contents are most immediately based on sense experience.

because aphasiacs or small children can think certain self-conscious thoughts that become unavailable to them if or when they (re)gain their linguistic capacity. This leaves unharmed a certain fallback principle, according to which the self-conscious thoughts of thinkers who *are* linguistically capable are linguistically expressible. Suppose the principle fails because even linguistically competent thinkers can, exceptionally, have certain extraordinary, ‘mystical’ thoughts that remain completely ineffable to them. This leaves unharmed a certain fallback principle, according to which self-conscious thoughts with contents of a kind that renders them apt to give the thinker reasons for beliefs about the relatively determinate layout of her empirical world, had by speakers, are linguistically expressible. Let’s say, then, that we cross out the Expression Principle in the Argument from Expression and replace it with the latter fallback principle, what we may call the Restricted Expression Principle. We may then add the very plausible lemma that the self-conscious thoughts that are the best candidates for being those in which perceptual concepts are exercised are ones with a content that renders them apt to give the thinker reasons for beliefs about the relatively determinate layout of her empirical world. Finally, we alter the conclusion to read that the perceptual concepts *possessed by speakers* are either lexical or demonstrative (or recognitional) concepts. If there were no errors in the original Argument from Expression apart from the supposed problems with the original Expression Principle, then these lightly revised premises would support this lightly restricted conclusion, about the perceptual concepts of *speakers*. Since the points made in the last section show, in effect, that the perceptual concepts of speakers do not in general take the form of either lexical, recognitional, or demonstrative concepts, this restricted conclusion is sufficient to refute, by reductio, conceptualism about the perceptual experiences of speakers. This would clearly be too

bad for conceptualists; they would not be interested in defending a claim about the perceptual experiences merely of non-speakers.

Ought we then to accept the Restricted Expression Principle? It seems pretty plausible that at least normally when a speaker has self-conscious thoughts with a content that renders them apt to give her reasons for entertaining certain beliefs about the relatively determinate layout of her empirical world, she can give at least a rough-and-ready linguistic expression to those thoughts. If this is so, but if the same does *not* go for the self-conscious thoughts in which perceptual concepts are used, the latter would to that extent be abnormal. What could account for that abnormality? It is hard to see what a *conceptualist* could do to answer that question. After all, the underlying idea in conceptualism is that the representational capacities on the exercise of which the contentfulness of perceptual experience depends are of a piece with the representational capacities that are exercised in normal, straightforward self-conscious thinking. I take it therefore that the conceptualist ought at the very least to accept the Restricted Expression Principle. Since the other steps of the Argument from Expression should be less controversial, I take that, given the Kantian Principle, the argument is sound, or at least ought to be by conceptualist's lights.

Another motivation for the claim that perceptual concepts are either lexical, recognitional, or demonstrative concepts is spelt out by what I call the Argument from Reflection. Its first three steps are shared with the Argument from Expression. Again, I will first state it in a rather bold version, and then suggest some possible refinements below.

The Argument from Reflection

- (1) Assume that conceptualism is true, and so that relative to an experience as of ...F..G..., there are certain concepts of F, of G, one needs to possess to have that experience. Call these 'perceptual concepts'.
- (2) In order for someone to possess a genuine concept, he must be able to exercise that concept in thinking self-conscious thoughts, that is, in thoughts that are available for immediate self-ascription (whether or not they are actually so ascribed). (The Kantian Principle)
- (3) So perceptual concepts must be exercisable in self-conscious thoughts. (From (1) and (2).)
- (4) Any self-conscious thought is available for introspective reflection, and by means of such reflection a subject can come to tell, *at least broadly*, what kinds of concepts are used in it. (**The Introspection Premise**)
- (5) So perceptual concepts must be exercisable in thoughts such that, by reflecting on those thoughts, a subject can come to tell, *at least broadly*, what kinds of concepts are used in it. (From (3) and (4).)
- (6) When competent subjects reflect to the best of their ability on the self-conscious thoughts that are the best candidates for being those in which perceptual concepts are used, they find no other candidates for being those concepts than lexical, recognitional, or demonstrative concepts. (**The Findings Premise**; motivated in terms of (5))
- (7) So perceptual concepts are either lexical, recognitional, or demonstrative concepts

Beside the Kantian Principle, the key steps here are the Introspection Premise, in (4), and the Findings Premise, in (6).

Now, someone might find the Introspection Premise implausibly strong. However, there are ways in which the Introspection Premise can fail that do not matter much to the spirit if not the letter of the Argument from Expression. Say, for example, the premise fails because there are various extraordinary self-conscious thoughts for which the subject cannot tell, however keenly she reflects on her thoughts, not even broadly, what kinds of concepts are used in them. Or suppose the premise fails because there are only certain highly reflective subjects, perhaps equipped with specialised theoretical training, able to tell, even broadly, what kinds of concepts are used in their self-conscious thoughts by reflecting upon them. If the premise fails for either of these reasons, it leaves unharmed a certain fallback principle, according to which philosophers (as an instance of highly reflective subjects, with specialised training) can tell at least broadly what kinds of concepts they use in thinking self-conscious thoughts of a certain mundane non-extraordinary kind. Suppose, then, we interchange this Restricted Introspection Premise for the original Introspection Premise in the Argument from Reflection. We also add a lemma to the effect that the self-conscious thoughts that are the best candidates for being those in which perceptual concepts are used are of the mundane non-extraordinary kind. If we take ‘competent subjects’ in the Findings Premise to be philosophers, the Restricted Introspection Premise along with the lemma motivates the Findings Premise no less well than the original Introspection Premise allows for.

Even if the original Introspection Premise, or the Restricted Introspection Premise along with the lemma, is granted, the Findings Premise does not of course logically follow; at best, it stands as a reasonable supposition in light of them.¹³ How might the motivation for it be spelled out a little more precisely? Well, one might think that if the perceptual concepts were of another kind, they would have turned up in the joint reflective efforts of philosophers by now. Tye suggests a related line of reply in the following passage.

Suppose that I am viewing a colored patch and that my visual experience conceptually represents this patch as red₂₅. Suppose further that my experience is not fleeting: I am staring at the patch for a considerable length of time. While my experience lasts, can I think to myself a thought which exercises this concept, for example, the thought that I am seeing something with shade red₂₅? It seems to me that the only thoughts I can form at such a time about red₂₅ have a demonstrative content. I can mentally “point” at the shade I am experiencing. I can think of it as that shade or that shade of red or perhaps just that. But, if my thoughts *seem* to me to have a demonstrative content, then, given that I have privileged access to the contents of my thoughts (that I can know via introspection alone what I am thinking) [MT adds in footnote *ad loc* ‘Assuming my faculty of introspection is working properly’], they do have such content. (Tye 2006a: 521-2)

Tye’s point here is that when subjects (presumably with a modicum of philosophical or psychological training) reflect on the thoughts that are the best candidates for being those in which a perceptual concept of red₂₅ is used, they get positive introspective

¹³ Since the role of the Introspection Premise (whether original or Restricted) is to make the Findings Premise seem a reasonable supposition, the qualification that reflecting subjects can tell ‘*at least broadly*’ what kinds of concepts are used in their self-conscious thoughts may be taken as emphatically as is consistent with the Introspection Premise non-trivially supporting the Findings Premise. To give an illustration of what this means, the Introspection Premise need not entail and ought not entail that reflective subjects can tell by reflecting on their thoughts whether demonstrative-concepts-as-construed-on-Evans’s (1982)-theory are used there, or whether demonstrative-concepts-as-construed-on-Pylyshyn’s (2003)-theory are used there. The Introspection Premise need not entail this, since no claim is made in the Findings Premise that subjects can make such theoretically deep judgements on the basis of introspective reflection. It is enough if the Introspection Premise entails that reflective subjects generally can tell, for example, whether or not a demonstrative, as opposed to a non-demonstrative-lexical concept, is used.

evidence that the relevant concept is demonstrative. Since there is no particular reason to think that introspection is misleading here, we ought to conclude that it indeed is.

If the Kantian Principle is true, the considerations spelt out by the Arguments from Expression and Reflection would, I reckon, give good reason for concluding that perceptual concepts are either lexical, or recognitional, or demonstrative. Given the conclusions in last section that subjects need not possess either lexical, recognitional, or demonstrative concepts of each of the features they experience things as having, this means that if the Kantian Principle is true, conceptualism is false. I conclude therefore that conceptualists are committed to giving up the Kantian Principle.

Conceptualists have not, to my knowledge, been at all attracted eager to do so. Yet it is not wholly clear to me why anyone should be so reluctant – why one should say, with McDowell (1994: 20), that it would be ‘fraudulent labelling’ to call anything not exercisable in thinking self-conscious thoughts a case of concept possession.¹⁴ The question becomes all the more pressing in light of the evident costs of accepting the Kantian Principle, outlined above. Given the conceptualisation constraint on belief and desires, it forces us to say that cats and dogs never want or believe anything. Given those constraints, and the plausible-sounding idea that tacit knowledge is knowledge, it forces us to give up the widely-held view that subjects may tacitly know things they are not capable of thinking self-conscious thoughts about. Perhaps solid reasons can be given for the principle, making these costs mandatory, yet, at least for the time being, it

¹⁴ It is no good to reply, as someone might be inclined to, that it is simply stipulative of a certain theoretical use of the term ‘concept’ that the Kantian Principle holds for it. Either this stipulation captures a ‘joint’ between two, interestingly different types of representational capacities or it does not. If it does not, why should the relevant theoretical use of the term ‘concept’ be an illuminating one? If it does, a non-stipulative reason in favour of the Kantian Principle could be given in terms of specification of the interesting differences between the two types of representational capacities.

is not clear to me why not giving up the Kantian Principle should not give at least temporary respite to the conceptualist.

3. Systematicity and its failures in perceptual experience

The last section argued that conceptualists are committed to abandon the Kantian Principle, but left open whether it is legitimate to do so. In leaving this open I did not intend to deny that there are costs to abandoning it (just as there are to affirming it). One such cost is that it leaves the conceptualist with one less option for a (perhaps partial) answer to a deep question that is bound to come up sooner or later, possibly with a critical intent, viz., what makes it right to say that perceptual concepts are really genuine, bona-fide *concepts*.

This deep question arises against the following background. Anyone who accepts that perceptual experience is representational or intentional accepts that in order to have a perceptual experience as of ...F---, a subject needs to have a perceptual-representational capacity to represent something as F. Conceptualists obviously accept this, but so do the vast majority of their contemporary critics. In the following, I will assume that this is common ground. The distinctive further claim of conceptualist is that the perceptual-representational capacity to represent something as F, needed in order to have an experience of something as F, constitute or include a *conceptual* representational capacity – a capacity that amounts to possessing a concept of F. The contrasting claim of the nonconceptualist is that the relevant perceptual-representational capacities are all *nonconceptual* representational capacities.

Now, if the nonconceptualist admits that the very idea of a conceptual representational capacity is coherent (as she surely must), she is committed to there being a non-arbitrary

distinction between conceptual and nonconceptual representational capacities. Similarly, if the conceptualist admits that the very idea of a nonconceptual representational capacity is coherent (as she probably ought to), she is subject to the same commitment. To get to the bottom of what is (or ought to be) at stake between the conceptualist and nonconceptualist we thus need to make clear what the presupposed principled distinction between conceptual and nonconceptual representational capacities is (or ought to be). This is clearly a deep and difficult question. However, in so far as the conceptualist and her opponent can agree that there are certain types of concepts, X, Y and Z, and that concepts of type X, Y and Z exhaust the candidates for being those possession of which is required to have a perceptual experience of something as F, the debate between them can proceed without going into that deep and difficult question. This fortunate state of affairs has often seemed to obtain in the debate, with lexical, recognitional and demonstrative concepts being the instances of X, Y and Z. However, if the conceptualist denies that concepts of these three types exhaust the candidate perceptual concepts, this fortunate state of harmony is disrupted. It becomes less legitimate to avoid the deep and difficult question, and correspondingly more legitimate for her opponent to press her for an answer.

Moreover, in giving up the Kantian Principle, the conceptualist has given up one resource for addressing that question. Unlike the nonconceptualist, she is not free to propose that conceptual representational capacities are while nonconceptual representational capacities are not exercisable in thinking self-conscious thoughts. If she is unable to give some other answer, she would clearly be left in an embarrassing position; McDowell's verdict of 'fraudulent labelling' might then be apposite.

What resources, then, are available to the conceptualist who has given up the Kantian Principle? One thought that may well appeal to this conceptualist is that concepthood perhaps essentially has something to do with what Evans called the Generality Constraint, or with the very closely related phenomenon that many others, most notably Fodor et al., have talked of as the systematicity of (certain forms of) cognition.¹⁵ Evans himself took conformity to the Generality Constraint to be a distinguishing mark of conceptual thought.

It is one of the fundamental differences between human thought and the information-processing that takes place in our brains that the Generality Constraint applies to the former but not to the latter. When we ascribe to the brain computations whereby it localizes the sounds we hear, we *ipso facto* ascribe to its representations of the speed of sound and the distance between the ears, without any commitment to the idea that it should be able to represent the speed of light or the distance between anything else. (Evans 1982: 104, nt. 22)

Fodor likewise takes concepthood to be characteristically linked with what he calls systematicity.

Qua constituents of thoughts, and of each other, concepts play a certain role in explaining why the propositional attitudes are productive and systematic. (...) Beliefs are *productive* in that there are infinitely many distinct ones that a person can entertain (given, of course, the usual abstraction from 'performance limitations'). Beliefs are *systematic* in that the ability to entertain any one of them implies the ability to entertain many others that are related to it in content. It appears, for example, to be conceptually possible that there should be a mind that is able to grasp the proposition that Mary

¹⁵ Evans illustrates the Generality Constraint thus: '[I]f a subject can be credited with the thought that *a* is *F*, then she must have the conceptual resources for entertaining the thought that *a* is *G*, for every property of being *G* of which she has a conception. This is the condition that I call "The Generality Constraint"' (1982: 104). For illustrations of Fodor's use of the term 'systematic', see the quote to follow presently in the text.

loves John but not able to grasp the proposition that John loves Mary. But, in point of empirical fact, it appears that there are no such minds. (Fodor 1998: 25-6)

There are interesting differences between how Evans seems to conceive of the Generality Constraint and its application to thought, and how Fodor conceives of systematicity and its application to conceptual cognition. For Evans, the Generality Constraint seems to have the status of an a priori truth, whereas for Fodor the systematicity of conceptual cognition is an a posteriori psychological law.¹⁶ This difference does not matter to our purposes; I will assume in what follows that the systematicity of conceptual cognition is *at least* a psychological law.¹⁷ Another apparent difference, more pertinent to us, is this. Evans, as we see, affirms that there are types of broadly cognitive representational capacities to the exercise of which the Generality Constraint does not apply, for example, the information-processing in the brain involved in auditory localisation of sounds. At least in some influential discussions, Fodor et al. voice much more scepticism over whether any cognitive representational capacities fail to be systematic. Responding to the suggestion that perhaps only the minds of verbal organisms are systematic, Fodor and Pylyshyn write:

Think what it would mean for this to be the case. It would have to be quite usual to find, for example, animals capable of representing the state of affairs *aRb*, but incapable of representing the state of affairs *bRa*. Such animals would be, as it were, *aRb* sighted but *bRa* blind since, presumably, the representational capacities of its mind affect not just what an organism can think, but also what it can perceive. In consequence, such animals would be able to learn to respond selectively to *aRb* situations but quite *unable* to learn to respond selectively to *bRa* situations. (So that, though you could teach the

¹⁶ Peacocke (1992: ch. 2) explicitly argues, in opposition to Fodor, that the Generality Constraint is an a priori principle. He attributes the same a priori view to Evans.

¹⁷ The law is only *ceteris paribus*, as virtually all laws outside fundamental physics are. We shall comment on the possible significance of this below.

creature to choose the picture with the square larger than the triangle, you couldn't for the life of you teach it to choose the picture with the triangle larger than the square.)

It is, to be sure, an empirical question whether the cognitive capacities of infraverbal organisms are often structured that way, but we're prepared to bet that they are not. (Fodor & Pylyshyn 1988: 27)

A little later, Fodor and Pylyshyn conclude, even more generally: 'That infraverbal cognition is pretty generally systematic seems, in short, to be about as secure as any empirical premise in this area can be' (1988: 28). Suppose they are right about this: that cognition, verbal and infraverbal, is generally systematic. Then systematicity could not be used, as the Generality Constrain can be used, on Evans's view, to draw a conceptual v. nonconceptual distinction *within* cognitive-representational capacities. In this situation, our current conceptualist would be faced with a choice. The 'hard-nosed' option would be to say that systematicity is still a feature that conceptual representational capacities have and nonconceptual representational capacities lack. This option is hard-nosed in so far as it commits one to denying, supposing that Fodor and Pylyshyn are right, that there are any cognitive nonconceptual representational capacities. Such capacities, to the extent that there are any, must be limited to the capacity of smoke to represent fire, of ice to represent temperatures below zero degrees centigrade, and the like. The 'soft-nosed' option would be to give up on systematicity as providing a principled mark of conceptual as opposed to nonconceptual representational capacities. In what follows, I shall assume that the conceptualist is well-advised to take the hard-nosed option. Indeed, I shall assume it is right to distinguish conceptual from nonconceptual representational capacities in terms, somehow, of the systematicity of the former. This has the consequence that if Fodor and Pylyshyn are right about the

systematic nature of cognitive capacities in general, conceptualism straightforwardly falls out as a truth about the special cognitive capacity of perceptual experience.¹⁸

Before we can assess whether Fodor and Pylyshyn are right about this we ought to make the way in which systematicity is supposed to be a mark of conceptual capacities a little clearer. Note that systematicity in the first instance is a feature of abilities to have thoughts, or, perhaps, of abilities to be in cognitive states of a certain kind. It is exemplified by its being a law that, if someone can think that he gives Mike Tyson a good bashing, he can think that Mike Tyson gives him a good bashing, and by its being a law that, if someone can think that Lisa embroiders and Tim hammers, he can think that Lisa hammers and Tim embroiders. Somewhat more generally, it is exemplified by its being a law that, if someone can think that a is F and that b is G, he can think that a is G and that b is F. This formulation does not capture the feature of systematicity in full generality, since systematicity has implications also for more structurally complex thoughts than such simple ‘atomic’ thoughts as those to the effect that a is F or that b is G. I will take this to be tacitly understood in what follows. Another drawback with the exemplifications of systematicity just given is that they only exemplify systematicity for abilities to have thoughts. Once we give up the Kantian Principle, we cannot take for granted that *all* conceptual capacities can be exercised in thinking thoughts. Given our purpose to formulate a mark of conceptual capacities, we ought, then, to characterise or illustrate systematicity in a way that makes it clear that we remain neutral on whether systematicity exclusively is a trait of the ability to have thoughts or whether it also applies to the ability to be in intentional or representational states of some other kind K

¹⁸ In saying that perceptual experience is a cognitive capacity, I do not intend to deny that it is not sensory, or has a distinctive phenomenology, only to register that it clearly falls under what is properly studied by cognitive science.

(where states of tacit knowledge, or perceptual experiences, may be such other instances of K).¹⁹ We may, then, take the following to be the salient characterisation, or illustration, of systematicity for our purposes:

Systematicity Intentional/representational states of kind *K* exhibit systematicity iff it is a law that, if one can represent that *a* is *F*, and represent that *b* is *G*, one can represent that *b* is *F*, and represent that *a* is *G*, where, in each case, the representational states are of the relevant kind *K* (and *mutatis mutandis* for more structurally complex representational states of kind *K*).

Given this characterisation of systematicity we can formulate the following principle for concepts:

The Systematicity Principle In order for someone to possess a concept, he must be able to use it having intentional states of some kind *K* that exhibit systematicity.

If this principle is true, but the corresponding does not hold for nonconceptual representational capacities, it can serve as giving at least part of the principled difference between conceptual and nonconceptual representational capacities.

Do, then, perceptual concepts meet the Systematicity Principle? That is to say, is it true that the concept of *F* that a subject needs to possess in order to have an experience of something as *F* (supposing conceptualism is true) meets the Systematicity Principle? When we raise that question, it quickly becomes apparent that the answer the

¹⁹ I treat 'intentional state' and 'representational state' as stylistic variants meaning the same.

conceptualist is in a position to return to it is closely linked with the fate of the Kantian Principle. Suppose that perceptual concepts essentially could be used in self-conscious thought. Since self-conscious thoughts plausibly are systematic, perceptual concepts would then plausibly meet the Systematicity Principle. It is thus only when the conceptualist gives up the Kantian Principle that the status of the Systematicity Principle, for her, becomes problematic.

Our current conceptualist cannot assume that perceptual concepts are useable in having thoughts to which the thinker has reflective access. Logically, she could suppose that perceptual concepts are useable in *subconscious* thoughts, such as states of tacit knowledge or information-processing states in the brain. But this supposition is problematic. First, as Evans notes, it is far from clear whether representational states of these kinds are systematic. Second, and more problematically, even if they are, the supposition that perceptual concepts are so usable seems to be a stab in the dark – seems to be a conjecture for which there is, to my knowledge, little empirical evidence. I set this suggestion aside, then, at least for now.

This leaves us with perceptual experiences, and, perhaps, sensory imaginations, as the only plausible candidate intentional states in which perceptual concepts are essentially useable. In turn, this means that whether perceptual concepts meet the Systematicity Principle hangs on whether perceptual experiences or sensory imaginations are intentional states exhibiting systematicity. I will now argue, pace Fodor and Pylyshyn, that they do not. There are various putative exceptions to systematicity in perceptual

experience (plausibly matched by corresponding exceptions for sensory imaginations, such as states of visual imagery).²⁰ I will note four kinds.

The ping-pong ball and the plate

The first type of case, I call the case of the ping-pong ball and the dinner plate. I can have a visual experience as of a dinner plate (oriented face-on) being behind a ping-pong ball centred on the plate (i.e. the centres of the ping-pong ball and the plate seem to line up in my line of sight). However, it seems that I cannot in the same sense have a visual experience as of a ping-pong ball being behind a plate (oriented face-on) that is centred on it. More generally, describing the experiences in simpler terms, we have the following exception to systematicity: I can have a visual experience as of a smaller, transparent or opaque thing being right in front of a larger opaque thing, where both things are right in front of me in the same line of sight. But I cannot have a visual experience as of a larger opaque thing being right in front of a smaller opaque or transparent thing, where both things are right in front of me in the same line of sight.

It has been objected to me that one can have an experience as of a ping-pong ball being behind the plate, for example when a sufficiently large part of the ping-pong ball sticks out from the side, or when the plate is transparent, or when the plate has somehow melted over the ping-pong ball, leaving a ping-pong-ball-shaped impression upon it. But these possibilities do not refute the exception to systematicity. What systematicity requires is that, given a possible experience A, then the experience B we get by recombining with respect to just *one* parameter in the specification of the content of A is

²⁰ I restrict myself to perceptual experience in the examples to follow (where this should be taken to include perceptual experience in cases of illusion and hallucination). I take the remarks to apply *pari passu* to (properly) sensory (as opposed to suppositional) imaginings (i.e. visual, auditory, or otherwise sensory imagery).

also possible. When I specified the first experience above, I was careful to say that the centre of the ping-pong ball was seen to line up with the centre of the plate. I might also have noted that the plate was seen as white (and so not as transparent), and as rigid (and so not melted), and so on. All these parameters are supposed to remain fixed in the second experience (the one I claim we cannot have) *except* that the ping-pong ball is seen as behind the plate. For this second experience, described in this way, no part of the ping-pong ball sticks out from behind the plate, the plate is not melted, and so on.²¹

Transparent white, reddish green, and other imperceptible colours

Another class of systematicity failures have to do with colour. Modern colour science, and indeed armchair reflection, has taught us that there are certain interesting limitations on the kinds of colour experiences we can have. Some of these give rise to exceptions to systematicity in perceptual experience. For example, Wittgenstein, in his *Remarks on Colour*, noted that there is no such thing as transparent white.²² This gives us the following exception. I can have an experience as of something being transparently blue, and an experience as of something being white, but I cannot have an experience as of something being transparently white.

Other possible failures of systematicity involving colour vision are these. Arguably, I can have a visual experience of something as redded yellow; for example, some or all experiences of orange may be appositely described in such terms. Arguably, I can have an experience of something as bluish green; for example, some or all experiences of

²¹ After writing the first version of this paper, Kent Johnson's (2004) intriguing article has come to my attention. Johnson notes a putative exception to systematicity in perception analogous to the one in the text: '[A]lthough I can see (imagine) a small black box inside of a large glass sphere, I do not think I can see (imagine) a large grass sphere inside a small black box.' (2004: 131) Johnson's paper is interesting notably in that it criticises the idea that systematicity applies even to thought and language.

²² Wittgenstein (1977)

cyan may be appositely described in such terms. But I cannot have an experience of something as redded green, or of something as bluish yellow.²³ The opponent-processing character of our systems of colour vision will not allow for such experiences.

The distant, biro-sized, biro-shaped thing

A third class of systematicity failures have to do with some inbuilt interconnections between the perception of size, shape and distance. The gist of these interconnections can be put as follows. Distant objects are not seen as having as detailed spatial outlines as close ones, unless the distant objects are seen as much larger than the close objects. Let's suppose that I'm looking at this pen under normal conditions, at arm's length. Let's describe this experience by saying that I see the pen as being biro-shaped, biro-sized, and at arm's length from me. It's important here that I take 'biro-shaped' and 'biro-sized' to express the highly fine-grained shape and size I see the pen to have under this close-up condition. Now consider a case where I see an object as being much further away, say about twenty meters. It is possible for me to see this object as biro-shaped, for example, if the thing is a huge, inflatable, biro-like object. But then I do not see it as being biro-sized but as being of a much larger size. Conversely, it is possible for me to see this somewhat distant object as biro-sized, for example, if the thing I'm looking at really is a biro. But then I don't see the object as biro-shaped, I see it only as roughly thin and straight, or something to that effect. So it seems I can't have an experience of something as relatively distant, biro-shaped and biro-sized. Visual experience will not allow for that.

²³ For more on this, see, e.g., Kaiser & Boynton (1996).

Metaphysical impossibilities

It seems plausible to suppose that, if someone can have an experience of something as a red square and an experience of something as a green triangle, he can have an experience of something as a green square, and an experience of something as a red triangle. Reflecting on the plausibility of such generalisations may inspire a measure of sympathy for Fodor and Pylyshyn's view that infraverbal cognition, including, presumably, perceptual cognition, is systematic. However, there are other experiences the specification of which we obtain by recombining on the specification of the contents of the experiences mentioned at the start of this paragraph, viz., the experience of something as a square triangle, and the experience of something as a red green thing. These do *not* seem to be experiences we can have. Subject to one or two special exceptions, it seems that we cannot have perceptual experiences, even in hallucination, as of metaphysical impossibilities.^{24, 25}

These four types of exceptions to systematicity in perceptual experience make it reasonable to conclude, I take it, that perceptual experiences are not intentional states of a kind K exhibiting systematicity. The same goes for states of sensory imagery, or properly sensory imagination. Since these states remain the only plausible candidate states in which perceptual concepts essentially are exercisable, for the conceptualist who has given up the Kantian Principle, it means that perceptual concepts flout the

²⁴ The plausibility of the corresponding claim about sensory imagination is, surely, part of what supports the idea that imaginability is a guide to possibility. See Gregory (2004) for a recent defence of the latter idea.

²⁵ The arguable exceptions to this rule include experiences of Penrose triangles or Escher drawings, and the waterfall illusion (see Crane 1988). However, these cases do not show that, for each impossible situation the specification of which is obtainable by recombining on the specification of the contents of possible experiences, there is a possible experience as of that impossible situation. See fn. 37 below for more on this.

Systematicity Principle, by the lights of this conceptualist. Another way of putting the same conclusion is this. We saw in the last two sections that the perceptual-representational capacities to represent F needed in order to have an experience of something as F are not in general exercisable in self-conscious thoughts. So, if the relevant perceptual-representational capacities amount to the possession of concepts, the Kantian Principle fails. What we have seen in this section is that the relevant perceptual-representational capacities are not exercisable in intentional states that exhibit systematicity. So if the perceptual-representational capacities in question amount to the possession of concepts, the Systematicity Principle fails.

4. Objections

In this section I address some objections to the claims and arguments of the last section.

The unperceivables are also unthinkable

When Evans formulates the Generality Constraint, he explicitly says that it should be taken subject to ‘a proviso about the categorical appropriateness of the predicates to the subjects’ (1982: 101, nt. 17). An example of what he seems to have in mind might be this. Even if you can think that the number 7 is prime and that Frank is happy, the truth of the Generality Constraint, or of systematicity, in no way should entail that you can think that the number 7 is happy, or that Frank is prime. This example involves cases of what we might call ‘text-book category mistakes’. Let’s suppose, then, that systematicity, even as formulated for the paradigm case of thought, ought explicitly to be hedged against text-book category mistakes. If this is so, one might think it ought also to be hedged against certain phenomena which, although not text-book examples of category mistakes, are somewhat related. The thought that something is transparently white, or that it is a square triangle, might be held to fall into this wider category.

Systematicity, even for thought, ought to be so formulated that although one can think that something is transparently blue and another thing white, this does not imply that one can think that something is transparently white. The same goes for such thoughts as that something is a square triangle.²⁶ Since this is so even at the level of *thought*, any more general formulation of systematicity, applicable also to perceptual experiences, will be similarly hedged. Thus at least many of our purported exceptions to systematicity in experience don't conflict with the law of systematicity properly formulated.

Now, even if this is so, it leaves standing the exceptions in the case of the ping-pong ball and the plate, and the distant biro-shaped, biro-sized thing. These cases involve no unthinkable thoughts on anyone's view. Further, even if systematicity for thought is to be hedged against text-book category mistakes (which is itself controversial),²⁷ it is a large claim to generalise this to the cases of transparent white and a square triangle. The words 'transparent' and 'white', or the concepts these express, clearly do not belong to different 'semantic categories' in the sense in which 'the number 7' and 'happy' may (i.e. in the sense that links up with '*category* mistake'). The generalisation would have to be justified on independent grounds. These grounds cannot be that the cases of transparent white, or of a square triangle, advert to metaphysical impossibilities, for we

²⁶ These restrictions on systematicity might be justified (and have in conversation been justified to me) by appeal to the idea that in order to think that something is so-and-so, one must know what it is for an arbitrary thing to be so-and-so. See Evans (1982: ch. 4) and Peacocke (1992: chs. 1-2) for defences of that idea. The claim, then, is that no one knows what it is for something to be transparently white, or for something to be a square triangle.

²⁷ See Magidor (2007) for a defence of the meaningfulness of category mistake sentences, and criticisms of restrictions on systematicity motivated by the idea that they are no meaningful.

think about (and indeed believe in) metaphysical impossibilities all the time.²⁸ Nor can the grounds be that these cases involve somehow logically incoherent suppositions. If *reductio ad absurdum* is to be possible, it must be possible for us to think, say, that P and not P. In light of this, and in the absence of a clear ground to the contrary, we ought to conclude that it indeed does follow from the systematicity of thought that one can think that something is transparently white, if one can think that something is transparently blue and another thing white, and *pari passu* for the thought that something is a square triangle.

The ceteris paribus character of the law of systematicity

Evans, as we saw, proposes that the Generality Constraint should be explicitly hedged against category mistakes. Whether or not this is a good idea, it is plausible that if the law of systematicity is considered an a posteriori psychological law about cognition, it is subject to the open-ended hedge of a *ceteris paribus* clause; after all, virtually all laws are plausibly so hedged, except, perhaps, for those of basic physics. The following toy model illustrates how a *ceteris paribus* clause might justifiably be invoked even for formulations of systematicity restricted to the case of thought. We suppose thinking is implemented on a simple language of thought type (LOT) system in the brain, where one thinks complex thoughts by shuffling around LOT tokens in the brain, and connecting them with each other in a designated thought-box. Suppose moreover that the LOT tokens might, by rare accident, become electrically charged, and that, in consequence, two similarly-charged tokens may repulse each other, making it impossible for them to co-habit in the thought-box. If the LOT tokens for ‘ice-cream’ and ‘good’ become both positively charged, then although our subject can think that something is good and that

²⁸ For example, I may believe that Peter is Frank’s son. Since he isn’t, what I believe is metaphysically impossible, on the plausible assumption of the essentiality of origin. Similarly, mistakes in calculation lead people to believe, say, that $46 + 237 = 273$.

ice-cream is something, he cannot think that ice-cream is good.²⁹ This is, of course, just one putative exception eligible to be swept under the *ceteris paribus* clause of the law of systematicity, restricted just to thinking. When we broaden our focus and formulate our law of systematicity in more general terms, so that it makes sense to apply it even to perceptual experiences, we would expect an even larger field of putative exceptions, apt to be swept under the *ceteris paribus* clause, to present themselves. Why should not all of our putative exceptions to systematicity in perceptual experience be of this kind? This question has bite since there admittedly are a number of more restricted generalisations that seem to instantiate systematicity in experience.³⁰

To answer this question we need to have some at least rough-and-ready idea of when it is appropriate to sweep a putative exception to a generalisation under its *ceteris paribus* clause, and when it is not. (Obviously, if it always is, the generalisation is trivial.) One plausible idea here is that it is not appropriate if the putative exception occurs, or would occur, under circumstances in which the idealisations relevant to the generalisation hold. Another, closely related, plausible idea is that this is appropriate just when the exception constitutes or reflects an independent interference on the system the relevant generalisation purports to describe.³¹ On both of these criteria, putative exceptions to laws of population genetics, say, that result from meteorite impacts, are rightfully covered by the *ceteris paribus* clause. It is part of the idealisations relative to these laws that significant meteorite impacts do not occur, and such impacts are paradigmatically independent interferences on the biological systems these laws purport to describe. The

²⁹ I vaguely remember reading this example in Peacocke, but I have been unable to locate its source.

³⁰ For example, as we noted above, it seems plausible that, if someone can have an experience of something as a red square and an experience of something as a green triangle, he can have an experience of something as a green square, and an experience of something as a red triangle.

³¹ See Pietrowski & Rey (1995) for a defence of broadly this view of *ceteris paribus* laws.

criteria similarly imply that our toy-model exception to systematicity for thought is rightfully covered. It is part of the idealisations of that law that the enabling or implementing mechanisms of cognition operate normally, yet the charging of LOT tokens was specified to be an abnormal happening. Similarly, these events whereby LOT tokens get charged in a way that relevantly affects their cognitive behaviour is aptly seen as an interference from a relatively independent physical domain upon the cognitive systems that the law of systematicity purports to describe. Neither can be said about the exceptions noted to systematicity in perceptual experiences. Laws that purport to state general features of cognitive states, cannot idealise away from lawful, non-accidental, cognitive features of those states. All of our exceptions clearly rest on lawful, non-accidental features of visual experience. Moreover, they are features of our systems for visual experience as such, not simply of an implementing domain. Non-accidental features of distance, depth, or colour-perception clearly are not independent interferences on the systems cognitive science purports to describe, as these clearly include perceptual systems. Our exceptions, then, are not rightfully covered by the *ceteris paribus* clause of the law of systematicity.

Misleading impression of structure in perceptual experience

When we attribute beliefs and desires, there is semantic structure in the part of the attributing sentence that specifies the content of the belief or the desire. We presume that this structure matches pretty well a certain structure in the representational capacities exercised in having the relevant belief or desire. For example, when we attribute to someone the belief that a sphere is transparently blue, we presume that he, in having this belief, exercises a representational capacity by which he represents something as transparent, *and* another, at least partly independent, capacity by which he

represents something as blue. His belief involves the joint exercise of several such, at least partly independent representational capacities. Or so we presume.

Now, when we attribute perceptual experiences, there is again semantic structure in the part of the attributing sentence that specifies the content of the experience. For example, when someone is said to experience a sphere as transparently blue, there is of course the same semantic structure in ‘transparently blue’ as there is when we attribute the belief. Now, someone might argue that there is, nevertheless, a disanalogy between perception and belief here, namely that the presumption of any corresponding structure in the representational capacities exercised in the perceptual experience is unfounded and probably false. Experiencing transparent blue might involve an as unitary representational capacity as experiencing blue does; we ought not to presume that it involves a capacity also used in experiencing opaque blue and *another* capacity also used in experiencing transparent red. Now, if all of this is so, someone might think it motivates an objection to the arguments of the last section, as follows. If these are what the underlying facts about the representational capacities involved in perception and thought are, then, since systematicity fundamentally describes the interaction of representational capacities, there is not so much as the expectation, on grounds of systematicity, that our ability to experience white and to experience transparent blue should come along with an ability to experience transparent white; hence our lack of the latter ability couldn’t possibly show anything about non-systematicity.

There is no need to disagree with anything before the last sentence here, which introduces the objection. On the contrary, the suggestion that it is a fallacy to infer structure in the representational capacities implicated in perception from semantic

structure in the phrases specifying its content could be *supported* on the basis of the arguments of the last section. What the objection gets wrong is that it puts the cart before the horse. Arguments over systematicity do not start by determining what the structure of the underlying representational capacities are, make hypotheses about how these underlying capacities may or may not combine with each other, and then infer, as implications of a principle of systematicity, that a creature with these capacities should be capable of having such-and-such intentional states.³² Quite the opposite, arguments over systematicity start by noting certain interesting patterns among the intentional states (or a certain kind) that a creature is capable of being in, where the patterns in question are broadly semantic. It is no doubt a delicate matter to make precise the sense in which they are semantic; still, what is clear is that the patterns manifest themselves in the specifications of the contents of the relevant intentional states. These broadly semantic patterns among capacities for intentional states are what arguments over systematicity affirm at the outset. The claim of systematicity should be identified with these affirmations.³³ Certainly, theorists who make these affirmations very often go on to infer, on their basis, hypotheses about the structure of the representational capacities involved.³⁴ But it is a mistake to think that we must know the truth of these hypotheses before we can assess claims of systematicity.

Other non-systematic but conceptualised states.

Another type of objection admits that perceptual experiences are not systematic whereas thoughts are. At any rate, it admits that there is a difference of kind between them in

³² Perhaps this gets the 'order of being' right. The claim is that it gets the 'order of discovery' wrong.

³³ In denying that perceptual experiences are systematic, we deny an affirmation of this kind about them.

³⁴ Fodor et al. most notoriously do this. But the same argumentative pattern can be discerned in Evans.

respect of systematicity. What it disputes is that this gives good reason to think that perceptual concepts are not bona fide concepts. For it argues that there is a range of other uncontroversially conceptualised states that fail to be systematic. If bona fide concepts are exercised in these non-systematic states, why should the concepts exercised in perceptual experience be any worse?

If there are factive mental states, such as knowledge, they saliently illustrate failure of systematicity. I can and do know that Las Vegas is east of Anchorage, but I cannot know that Anchorage is east of Las Vegas, because it isn't. Arguably, there are putative exceptions to systematicity also for desires. Someone might desire that he be better than everyone else, yet claim that he is wholly incapable of desiring that everyone else be better than him. Some thoughts may be just too awful for anyone to desire them true. Similarly, some thoughts may be just too obviously false for anyone to believe them, perhaps outright contradictions are a case in point. That would imply exceptions to systematicity for belief. Yet beliefs and desires are paradigm conceptualised states, and knowledge certainly conceptualised if a mental state at all.

Now, the special systematicity-exceptions for factive states that arise from the fact that they cannot have false contents seem to me to be beside the point. There is no psychological explanation to be given of why we cannot know that Anchorage is east of Las Vegas. This contrasts with all the other putative exceptions to systematicity we have listed, where there is a psychological story to be told of why we cannot be in the relevant intentional state (if indeed we cannot). For this reason, if not for others, these systematicity-exceptions seem to be sufficiently different from the rest to set them aside.

Turning to beliefs, it is not clear that there are thoughts too obviously false for anyone to believe. Dialethic logicians claim that some straightforward contradictions are true.³⁵ Presumably they believe them. Moreover, they believe them not as a result of some wild freak of psychology, but as a result of reasoning. They may indeed have made various mistakes in their reasoning, but as we very often do, this does not show that they have not arrived at them by an eminently normal way of belief formation. If a normal thinker had been presented with their reasoning, had the time and patience to follow it in detail, and, for good measure, been given some antecedent strong reason to think that the reasoning ought to be compelling, then it is not implausible that he would have come to believe the contradictions in question too. If this is how it is, we cannot say he was incapable of believing the contradiction. Now, if arguments can be given in favour of straightforward contradictions that seem convincing even to skilled reasoners, is it reasonable to hypothesise that there is no proposition we can grasp for which such arguments could not be given. But if this is so there is no proposition we can grasp we are incapable of believing.³⁶

Parallel points extend to desires. Plenty of things are absolutely awful, and ought, perhaps, to seem absolutely and overwhelmingly awful to anyone. Yet arguments could be given that would make something objectively awful seem not so bad after all. The arguments may be fallacious, but as arguments often are, this does not mean that desires formed as a result of being swayed by them are not arrived at in a perfectly normal way. For example, someone might be convinced by a fallacious argument that if, but only if,

³⁵ See Priest (1987).

³⁶ I am here indebted to some analogous arguments in Williamson (2006), directed against the claim that there are certain beliefs one must have in order to possess certain concepts.

everyone else were better than him, some absolutely wonderful consequences would ensure. Being so convinced, it would be perfectly natural for him to desire that everyone else be better than him. Since this is what he would desire, if he were swayed by the argument (as he would be, had he been presented with it, had the time and patience to follow it in detail, etc.), we cannot claim that he was incapable of desiring it in the first place. Thus it is doubtful whether there are things too obviously overwhelmingly awful for anyone to desire.³⁷

Even if this is wrong, and beliefs and desires are indeed not systematic, the current objection misses an important point. It is uncontroversial that someone can believe, or desire, or know that P, only if they can think that P. Indeed, in believing, or desiring, or knowing that P, they entertain the thought that P (although often not occurrently). The concepts exercised in belief etc. are ipso facto exercised in entertaining thoughts. Since the current objection in no way disputes that entertaining thoughts is systematic it in no way disputes that the concepts used in having beliefs meet the Systematicity Principle. This contrasts with the case of perceptual experience. In giving up the Kantian

³⁷ Perhaps someone will raise the following objection at this point. Above, I admitted that there are (or may be) exceptional cases of experiences as of impossible situations, for example in the waterfall illusion. I denied, however, that these examples give good reason to think that perceptual experiences as of impossible situations of various kinds more generally are possible. In this section, I have just argued that belief in putatively obvious impossibilities, or desire for putatively overwhelming calamities, is possible by giving one or two examples of how this might come about. The question is why we should think these examples give better reason for generalisation about beliefs and desires than the corresponding examples for experiences (the waterfall illusion, Penrose triangles, etc.) do for generalisation about them. Am I not guilty of double standards here? I reply that I am not. The reason is that the examples in the doxastic/conative case trade on a mechanism that is un-encapsulated and domain neutral, viz. the capacity for persuasion. (See Fodor (1983) for the notions of (un-)encapsulation and domain neutrality invoked here.) For this reason, the mechanism is generalisable to belief in putatively obvious impossibilities, or desire for putatively overwhelming calamities, of other kinds. The specific perceptual mechanisms that underlie the experiences as of impossibilities are encapsulated and domain specific, and thus do not in the same way underwrite generalisation to experiences of impossible situations of other kinds.

Principle, the conceptualist denies that perceptual concepts must be exercisable in self-conscious thoughts. Since it is doubtful to say the least that they must be exercisable in subconscious thoughts, the only states they plausibly essentially are exercisable in are perceptual experiences (and perhaps sensory imaginings). For that reason, the failure of perceptual experience to be systematic shows something important, which the failure of belief (if indeed it fails) to be systematic does not, viz. that the representational capacities used in the relevant state fail the Systematicity Principle.

5. Conclusion

My aim in this paper has been to make clearer what the commitments of conceptualism are. In particular, I have been arguing that these include the commitment to give up two principles about concepts that most conceptualists, as well as many of their opponents, have wanted to affirm: the Kantian and the Systematicity Principles. I have not argued that either of these principles is non-negotiable, and so that conceptualism is refuted. Indeed, I have underlined the significant costs of accepting the Kantian Principle, and my failure to have found a clear and convincing reason for it.

What about the Systematicity Principle? Now, if self-conscious thought is systematic, as it has widely been held to be, and as we may assume *pro tem*, the Kantian Principle entails the Systematicity Principle. If, moreover, self-conscious thought is the *only* kind of intentional state exhibiting systematicity, the Systematicity Principle entails the Kantian Principle, and is just as costly as it. If, however, self-conscious thought is not the only kind of systematic intentional state, the Systematicity Principle does not entail the Kantian Principle, and is easier to accept than the latter. Since the claim that self-conscious thought is a systematic intentional state is less controversial than the claim that it is the *only* kind such a state, the Systematicity Principle ought, correlatively, to be

less controversial than the Kantian Principle. Giving it up is thus comes, dialectically, at a bigger price for the conceptualist.

Indeed, if thought is systematic (whether generally, or only in the shape of self-conscious thought), the conceptualist will, in giving up the Systematicity Principle, face an embarrassing question. If thought is systematic while perceptual experience (as we have seen) isn't, and if the representational capacities used in the latter are not in general exercisable in the former (as we have argued to be plausible), then it is reasonable to infer some interesting contrast between the representational capacities used in thought and those used in perceptual experience. Why shouldn't this contrast constitute (part of) the sought-after principled distinction between conceptual and nonconceptual representational capacities? If this contrast doesn't constitute it, what does?

As far as I can see, the only way the conceptualist can make this question less embarrassing, is by taking the radical tack of arguing that not even thought exhibits systematicity.³⁸ If that is so, the Systematicity Principle does not fail simply for some pretty peripheral reasons over perceptual concepts but was ever a non-starter. To establish this is of course not to find an alternative principled difference between conceptual and nonconceptual representational capacities. It would mean, though, that the problem of finding one is everyone's problem, not especially embarrassing for the conceptualist. Judging whether the radical tack could succeed is a task for another occasion; what is clear is that upholding conceptualism involves radical commitments.

³⁸ See Johnson (2004) for some recent interesting arguments to this effect.

A Problem about the Compatibility of Content Conceptualism with a certain Thesis of Diaphanousness

0. Introduction

The great majority of contemporary philosophers of perception assume that perceptual experiences, like beliefs or thoughts, have representational content.¹ A not insignificant subset of these philosophers hold, moreover, that that the representational content of a perceptual experience is identical to the content of a possible belief or thought.² This claim has come to be referred to as ‘content conceptualism’, and I will adopt that label. Both parts of the label call for a brief remark. The reason for the ‘content’ part of the label is to distinguish the claim of content conceptualism from the different (albeit perhaps related) claim of *state* conceptualism, according to which perceptual experiences have their content in a conceptualised way, i.e. are such that a subject can

¹ See Essay 1 above for a discussion of this claim, and defence of it against Charles Travis’s (2004) recent objections.

² Some of those who agree with the spirit of the claim that the content of a perceptual experience is identical to the content of a possible belief or thought may worry that the representational contents of some or all perceptual experiences are so rich as to render it psychologically impossible for anyone like us to embrace them within a single belief or thought. If this is their only worry about the target claim, then many of them would presumably be happy to agree that, for any perceptual experience, there is a partitioning of its representational content into parts each of which is the content of a possible belief or thought. One way to respond to the richness worry is thus to weaken the formulation of the target claim as follows. Any perceptual experience has a representational content such that there is a partitioning of its content into parts each of which is the content of a possible belief or thought. I will however operate on the presumption that those that are happy to accept this weaker claim ought to be happy to accept the target claim in the text as a legitimate idealisation. The idealisation in question can be seen as case of abstracting away from performance limitations in psychological theorising. The experiences we will be considering will moreover be of a relatively simple kind (e.g. experiences of uniformly coloured walls) for which the idealisation seems especially justified.

have a perceptual experience with the content that ...F... only if she has a concept of F-ness.³ State conceptualism will at best indirectly be in question in this paper, to the extent that it is connected with content conceptualism.

The motivation behind the ‘conceptualism’ part of the label is the idea that those who subscribe to the view subscribe to the claim that perceptual experiences (exclusively) have conceptual representational content (for short, conceptual content).⁴ Now, there is a thin sense of ‘conceptual content’ on which it is indeed a fair characterisation of the view to say that those who subscribe to it hold that perceptual experiences have conceptual content: viz. the sense of ‘conceptual content’ on which it just means ‘content of the kind that thoughts and beliefs have’. However, there are more loaded senses of ‘conceptual content’ on which this is not a fair characterisation of the view. An example of such a more loaded sense is that on which ‘conceptual content’ means content composed out of concepts. It is true that some content conceptualists hold that the contents of both experiences and thoughts take the form of fine-grained, neo-Fregean Thoughts.^{5, 6} The modes of presentations out of which such neo-Fregean

³ The distinction between content and state conceptualism is made explicit in Byrne (2005). In drawing the distinction, Byrne develops related suggestions in Stalnaker (1998) and Heck (2000).

⁴ A distinction is sometimes drawn between views on which perceptual experiences exclusively have conceptual content, on which they exclusively have nonconceptual content, and on which they have both conceptual and nonconceptual content. McDowell (1994, 1998a) and Brewer (1999) defend a view described in the first terms, Tye (1995, 2000) one in the second terms, and Peacocke (1992) one in the third. Now, our formulation of the thesis of content conceptualism uses neither of the terms ‘conceptual content’ or ‘nonconceptual content’. But a corresponding threefold distinction could be drawn within the framework of our formulation, between the claim (i) that perceptual experiences exclusively have content which is the content of a possible thought; (ii) that they exclusively have content that is not the content of a possible thought; and (iii) that they both have content that is the content of a possible thought, and content that is not. For our purposes, I shall identify content conceptualism with the first of these claims. I will let this restriction be tacit in what follows.

⁵ For the record, ‘thought’ is used in this paper not for the abstract things that are the things we think but for states of thinking. ‘Having a thought’ might be read as expressing a determinable whose

contents are built qualify as concepts, on one good sense of ‘concept’.⁷ However, other philosophers who endorse content conceptualism hold that the contents of both experiences and thoughts take the form of neo-Russellian propositions, or of functions from possible worlds to truth-values, or of possible individuals.⁸ Someone who combines content conceptualism with one of the three latter views is not committed to holding that perceptual experiences (or indeed beliefs) have contents composed out of concepts. Thus content conceptualists are committed to the claim that perceptual experiences have conceptual contents only on a certain thin sense of ‘conceptual content’.⁹

Many contemporary philosophers of perception also believe that perceptual – in particular, visual – experience is diaphanous or transparent.¹⁰ The idea goes back at least to Moore who observed that ‘when we try to introspect the sensation of blue, all we can see is the blue; the other element is as it were diaphanous.’ (1903: 25) The claim that visual experience is diaphanous is, then, a claim about what our experiences seem to us to be like, when we turn attention to our own experiences. The idea is that trying to ‘get at’ the phenomenal character of our own experiences is a bit like trying to

determinates include having a belief, entertaining a supposition, considering a hypothesis, etc. When used as a designation for the contents of thoughts and beliefs, as in ‘Fregean Thoughts’, I will capitalize it.

⁶ This is the view of McDowell (1994, 1998a) and Brewer (1999, 2005).

⁷ See Peacocke (1992) for a use of ‘concept’ in this sense.

⁸ For the first view, see Thau (2002) and Byrne (2005); for the second, Stalnaker (1998); for the third, Lewis (1983).

⁹ Stalnaker (1998) is an example of a writer operating with a more loaded sense of ‘conceptual content’. He argues that neither perceptual experiences nor beliefs have conceptual content, in this loaded sense. On the noted thin sense in the text, beliefs of course trivially have conceptual content.

¹⁰ I will restrict myself to the visual case in what follows, and talk in terms of ‘diaphanousness’ rather than ‘transparency’.

visually attend to the empty air a little in front of a surface one is looking at.¹¹ If I try the latter, my attention as it were ‘falls through’ the air and onto the surface in question. Analogously, when I try to get at the phenomenal character of experience it seems I can’t help attending to the apparent objects of the experience and to the properties they appear to have.¹² As Gilbert Harman wrote:

Look at a tree and try to turn your attention to intrinsic features of your visual experience. I suggest that the only features there to turn your attention to will be features of the tree. (1990: 667)

The minimal claim which simple phenomenological reflection along these lines seems to sustain is that in reflecting on the phenomenal character of our experiences we can’t help attending to the apparent objects of experience and their apparent properties. Many would agree (Harman certainly seems to be one of them) that a slightly stronger claim also is sustained by these reflections. This claim is what I will single out for our purposes as *the* thesis of diaphanousness (although there are certainly other theses in the nexus of diaphanousness considerations that equally could have been given that label):

¹¹ ‘Phenomenal character’ is here used a neutral term for what is it like for the subject to be in a certain state of mind.

¹² The phrase ‘apparent objects of experience’ is not intended to denote a special kind of objects. It simply denotes objects appearing to one in a given experience. Apparent objects of experience may turn out to be always real existent objects, sometimes merely intentional objects (as argued by Harman (1990) among others), or something else. I stay neutral on this question here.

(D) Thesis of Diaphanousness In introspection, one is or becomes aware of the phenomenal character of visual experience only *by* attending to apparent objects of experience and properties they appear to have.¹³

In this paper, I will raise a problem about the compatibility of content conceptualism with (D). The problem arises against the background of a certain plausible background assumption, which ought to be made explicit. The assumption concerns the relation between being in a state of mind with the sort of representational content that visual experiences have, on the hand, and being in a state of mind with the sort of phenomenal character that visual experiences have, on the other. It says that the former is insufficient to the latter. Consider, as an arbitrary example, the visual experience I now have as I look out of the window, in which various trees and buildings appear to me as having various shapes and colours etc. Uncontroversially, this experience has a certain phenomenal character. If representationalism is right, the experience also has a certain representational content, call it *p*. Now the question is this. Is being in a mental state with the content *p* (at a given time) a sufficient condition for being (at that time) in a mental state with a phenomenal character of a certain sort, viz. a phenomenal character typical of visual experiences with the content *p*? According to the relevant background assumption it isn't. I will refer to this assumption as 'Insufficiency'.

Insufficiency would seem to be a highly reasonable claim from the point of view of content conceptualists. They hold that there is a possible thought or belief with the same content *p* as my visual experience. If I have enough concepts, and don't suffer from any

¹³ The formulation of this thesis is inspired by theses (2) and (2c) in Stoljar (2004). (His thesis (2) is what I will label thesis (S) below.) My discussion of the diaphanousness considerations is much indebted to Stoljar's paper.

relevant performance limitations, I can have this thought or belief *p*. It seems highly plausible that, if I can have this belief at all, I can have it at a moment in time when I don't have any visual experiences, or entertain in any visual imagery, with a phenomenal character even remotely like that of my visual experience as I look out of the window. So it seems highly plausible that Insufficiency is true.¹⁴

At the same time, the plausibility of Insufficiency does not presume content conceptualism. Not only is it consistent with content *non*conceptualism, but some fairly natural elaborations of that view will be committed to it. Michael Tye (2000), for example, defends a content nonconceptualism on which perceptual experiences have neo-Russellian propositions as contents while thoughts have Fregean Thoughts as contents. Someone who shares this view may think that even if thoughts cannot have the same contents as visual experiences, there are other kinds of psychological states that (i) can have such contents and (ii) one can be in without being in any state with a visual phenomenal character; for example, subconscious representational states implicated in visual processing might be thought to meet these conditions.¹⁵ This fairly natural elaboration of a content nonconceptualist view is committed to Insufficiency.

1. The problem

Consider a simple visual experience, say one as of a blue wall in front of one. Let's call it 'Wall-Experience'. The content conceptualist holds that there is a possible thought with the same content as Wall-Experience. Let's call this thought 'Wall-Thought'. If we

¹⁴ Interestingly, not all content conceptualists endorse Insufficiency. Thau (2002) quite explicitly denies it. We shall return to this below.

¹⁵ Tye certainly thinks there are subconscious states with representational content implicated in visual processing. However, I cannot find a passage where he explicitly commits himself to the claim that such states might have the same contents as conscious visual experiences. The fairly natural elaboration of his view may thus not be one he would himself want to endorse.

assume Insufficiency, someone who has Wall-Thought need not be in a state of mind with a visual phenomenal character remotely like that of Wall-Experience (his eyes may be closed). Now consider a subject who, up until t , has Wall-Thought but no relevant visual experience, but who at t gets to have Wall-Experience. Such a subject at t undergoes a significant shift in the overall phenomenal character of her state of mind. Such significant shifts may be and typically are introspectively detectible. They may be and typically are things we can know by introspection.¹⁶

Yet this conclusion conflicts with the upshot of the following argument from diaphanousness:

- (1) Thesis of Diaphanousness, (D).
- (2) If (D), then, if entering into a certain visual state of mind involves no change in representational content, introspection will register no change in phenomenal character associated with the transition into the relevant state of mind.
[Assumption]
- (3) So, if entering into a certain visual state of mind involves no change in representational content, introspection will register no change in phenomenal character associated with the transition into the relevant state of mind. [MP from (1) and (2).]
- (4) Having Wall-Experience from t when one had Wall-Thought up until t does not involve any change in representational content. [Consequence of content conceptualism]

¹⁶ Block (1998) and Byrne (2001) argue that it is a conceptual truth that large changes in phenomenal character are introspectible. Even if this is not a conceptual truth, it surely seems to be a truth.

- (5) So, introspection will register no change in phenomenal character associated with the transition to having Wall-Experience at t when one has Wall-Thought up until t . [MP from (3) and (4).]

As we saw above, the negation of (5) is extremely plausible if content conceptualism and Insufficiency are true. So it is extremely plausible that (5) is inconsistent with the conjunction of content conceptualism and Insufficiency. Now, the argument is clearly valid, and claim (4) clearly follows from content conceptualism. Thus, if claim (2) is true, the argument shows that our Thesis of Diaphanousness, (D), is incompatible with the conjunction of content conceptualism and Insufficiency. Since Insufficiency is highly plausible if content conceptualism is, the argument plausibly shows, if (2) is true, that (D) is incompatible with content conceptualism.

The motivation for (2) is as follows. According to (D), in introspection, one is or becomes aware of the phenomenal character of visual experience only *by* attending to apparent objects of experience and properties they appear to have. Now any representationalist view of experience assumes that these apparent objects of experience, and the properties they appear to have, are precisely the objects and properties that figure in, or are determined by, the representational content of the experience.¹⁷ If the representationalist is to be motivated in her insistence that visual experiences have representational content, the notion of representational content must account for some salient feature of experience. If the notion accounts for *anything* about visual experiences, then surely it must account for the fact that objects appear to us in

¹⁷ The parentheses are included to stay neutral between a neo-Russellian view of representational content, on which contents are built up from elements on the level of reference, and a neo-Fregean view, on which contents are individuated in terms of modes of presentation of objects and properties.

visual experience, and appear to us to be certain ways – to have certain properties. The way it does so is by taking these objects and properties to be determined by the representational content of the experience. Since the content conceptualist is a representationalist, she is committed to agree with this.

If this point is agreed, we ought to accept that (D) has the implication that, in introspection, one is or becomes aware of the phenomenal character of a visual experience by attending to its representational content. This consequence is immediate if the representational content simply is built out of the objects that appear to us and the properties they appear to have, as it does on a neo-Russellian view. Yet even if the objects and properties in question are only determined from, and do not figure as constituents in, the representational content, it would be natural to assume that introspective access to these objects and properties is via the feature of experience that determines them, viz. the content. Moreover, the representational content determines the objects and properties in question in a way that is epistemically unproblematic to introspection. It is not as if the claim that we have introspective access to the objects that appear to us and the properties they appear to have via introspective access to the representational content of the experience makes introspective access to the objects and properties in question epistemically problematic in any sense.¹⁸ Therefore, the assumption that introspective access to these objects and properties is via introspective awareness of the content does not make introspection any less epistemically powerful than it would be if we somehow had ‘direct’ access in introspection to the object that appear to us and the properties they appear to have, i.e. if we had access otherwise than

¹⁸ The relation between the representational content of an experience and the objects and properties it determines is thus not to be confused with the relationship between, on the one hand, the sense-data that a representative realist posits as direct objects of perception and, on the other, the physical objects that these sense-data in some sense ‘belong to’, or are ‘representatives of’.

via attention to the content. In assuming that (D), for the representationalist, has the consequence that one is or becomes aware of the phenomenal character of a visual experience by attending to its representational content, we are therefore not illegitimately constraining the epistemic power of introspection.

Let's assume, then, that there is a change in phenomenal character of your state of mind at t , in that you at t come to have visual experience where you had none before, and that you become aware introspectively of this change in phenomenal character. It follows from (D), we have argued, that you become aware of the character of the state of mind you are in at t by attending to the representational content of your state of mind. Yet if there is no change in representational content involved in the transition into a visual experience at t , then how, by attending to the representational content of one's state of mind, could one detect a change in phenomenal character? If one succeeded in detecting such a change, it would seem that one would have to detect it by other means. But, given what we have just argued, (D) has the consequence that it is *only* by means of attending to representational content that one gets at the phenomenal character of visual experience. So it seems that the introspectively detectible change in phenomenal character of your state of mind at t , when you came to have a visual experience, must have involved a shift in the representational content. Contrapositively, if entering into having the visual experience involved no change in representational content, introspection would register no change in phenomenal character associated with the transition. If we conditionalise on (D), we thus get claim (2).

The conceptualist, being a representationalist, ought then to accept (2). Since she is also committed to step (4), she would be committed to (5) if she accepts (D). Since (5) is

incompatible with the conjunction of content conceptualism and Insufficiency, and since Insufficiency is a highly plausible claim for a conceptualist, it seems she is committed to deny (D).

2 Two objections to (2), and replies

I will now consider two objections to (2). They both trade on the fact that (D) only says how we get at the character of visual experiences; it does not say how we get at the phenomenal character (if any) of our thoughts.¹⁹ Thus, perhaps the way in which we get at the phenomenal character of thoughts allows for the possibility of registering a change in character even if there is no change in content.

The first objection spells out this worry in terms of the notion of ‘attention’ figuring in (D). We need to recognise a distinction between visual, or more generally perceptual, attention and what we might call ‘cognitive’ attention. The former is exemplified when we ‘experientially highlight’ some object or feature in our field of view, the latter when we intellectually concentrate on some particular problem, such as when we turn our attention to the validity of some inference.²⁰ Now it is arguable that the form of attention relevant to (D) is perceptual attention. The form of attention relevant to getting at the character of thoughts, by contrast, cannot be perceptual but has to be cognitive. Thus even if we get at the character of thoughts by attending to content, as we do for the case of visual experience according to (D), the different *form* of attention in the two cases may make for an introspectible difference.

¹⁹ I let the qualification ‘(if any)’ be tacit henceforth.

²⁰ See Martin (1997) for a discussion of this contrast between perceptual and ‘cognitive’ attention.

I reply that a change in the means of access must not be confused with means of access to a change. If we got at the character of perception in part by higher-order monitoring of the style of attention used to get at perceptual content, then the change of means might make the change in character introspectible. But this is not what (D) states: according to (D) we get at visual character by attending to how things appear to us, not by attending to how we thereby attend.

The second objection raises the possibility of us becoming aware of the character of thoughts not only by attending to their representational content, but also by responding to some other feature of thoughts. By responding to this further feature we become aware that thoughts have a (let's say) 'notional' – as opposed to perceptual – character. The suggestion is that this explains how we can introspectively register a change in character when we move from Wall-Thought to Wall-Experience: we register first a notional character and then a perceptual character.

I do not want to deny the possibility that we respond to some feature of thoughts beyond content; however I dispute whether it can explain the introspective registration of a change in character. The possibility that we respond to such a further feature could explain a change *in what we register about* the phenomenal character of our state of mind, in the following way: in reflecting on Wall-Thought we register a notional character by responding to some feature beyond content; in reflecting on Wall-Experience we do *not* register such a notional character, for we do not introspectively respond to such a further feature. However, the absence of a registration of a notional character in Wall-Experience must not be confused with the registration of the absence

of notional character. To explain the introspective registration – i.e. the introspective knowledge – of a change, we need the latter, not merely the former.

I grant that it is sufficient to register the absence of notional character to register the presence of something incompatible with notional character. The most likely candidate for something incompatible with notional character is of course perceptual character. Thus if one registers the presence of perceptual character in reflecting on Wall-Experience, I grant that one has what it takes to register a change. Yet how are we to explain one's registering a perceptual character? Assuming (D), the registration must be by virtue of attending to the representational content. However, suppose this explanation is sufficient. Then since we can attend to the same representational content in Wall-Thought, we should register perceptual character in Wall-Thought too. But this conflicts with the assumption that we register notional character there, and that registering perceptual character is registering something incompatible with notional character. It seems, then, that (D) really has the consequence that (2) ascribes to it.

3. Options for the content conceptualist

If (2) is true, what other options does the content conceptualist have? As far as I can see, the only remaining options are to deny Insufficiency or the Thesis of Diaphanousness, (D). I will consider these in turn.

Denying Insufficiency

Our case for the incompatibility of content conceptualism and (D) took the form of an argument, on the one hand, from content conceptualism and Insufficiency to the conclusion that (5) is false, and another argument, from content conceptualism and (D), to the conclusion that (5) is true. However, if Insufficiency is false – i.e. if Sufficiency

is true – then the result that (5) is true need not conflict with content conceptualism. Consider the relationship between Wall-Thought and Wall-Experience. The former has the same content as the latter. Sufficiency thus entails that if someone has Wall-Thought, they must be in a mental state having a phenomenal character typical of visual experiences having the same content as Wall-Experience. What phenomenal character is this? It is independently plausible to assume that if visual experience A and visual experience B have the same content, they have the same phenomenal character.²¹ So the phenomenal character typical of visual experiences having the same content as Wall-Experience is just the phenomenal character of Wall-Experience. This means that, given Sufficiency, someone who has Wall-Thought must be in a state of mind having the same phenomenal character as Wall-Experience. If someone up until t has Wall-Thought and has Wall-Experience from t (and has no other mental states with any relevant visual phenomenal character around t , as we may assume is the case), then *no* change in visual phenomenal character takes place at t . So the claim that our subject cannot introspectively detect such a change at t , as (5) affirms, is just the conclusion we should expect.²²

²¹ See Byrne (2001) for sustained discussion and defence of this claim. The claim in the text instantiates what Byrne (2001) calls a claim of restricted, intra-modal intentionalism.

²² Someone might want to argue as follows at this point: What is it for something to be a certain visual experience – to be a certain visually experiential state of mind? It is for it to have a certain visual phenomenal character. Thus, if a state of mind has precisely the phenomenal character of Wall-Experience, that state of mind *just is* Wall-Experience; which being so, it follows from the conjunction of content conceptualism and Sufficiency that it is impossible to have Wall-Thought at t without having Wall-Experience at t . I don't have to disagree with anything in this remark; in fact, it can be seen as an alternative way of putting the point in the text. The person who makes this argument agrees that, given content conceptualism and Sufficiency, a situation described merely as being one in which the subject has Wall-Thought until t and Wall-Experience from t and in which the subject has no other visually relevant states of mind, *cannot* be one in which there is a change in the visual phenomenal character of the subject's state of mind at t .

An example of a content conceptualist who quite clearly is committed to Sufficiency is Michael Thau (2002). He writes:

[P]erceptions are distinguished by what properties they represent objects as having and that's all there is to the phenomenology of perception. But, given this, what distinguishes perception from belief is that these properties aren't available in belief except by demonstrating them *while* perceiving them. (Thau 2002: 223, emphasis added)

Thau argues that the contents of both perceptual experiences and beliefs are 'Millian' or neo-Russellian. He does not deny the possibility to have thoughts about the objects one experiences to the effect that these objects have just the properties one experiences them to have. That is to say, he does not deny that this is possible *provided* a stringent condition is met: viz. that one at the same time has experience of these objects as having these properties. Given his neo-Russellian view of content, it follows that it is possible to have thoughts with the same content as one's experience provided that this stringent condition is met. But he is committed to hold that this stringent condition *must* be met for it to be possible to have such an overlapping thought. For he argues on the basis of considerations of diaphanousness that the perceptual phenomenal character of one's state of mind when one has a perceptual experience is determined *simply* by the fact that the representational content of the perceptual experience is what it is (i.e., given his neo-Russellianism, *simply* by 'what properties they represent objects as having'). So if someone has a thought with the same content *p* as a visual experience, she must be in a state of mind with the visual phenomenal character that is typical of visual experiences with that content *p*. That is just to say that Sufficiency holds.

Thau's combination of content conceptualism, neo-Russellianism, and Sufficiency is highly implausible. If one can have beliefs about the objects as one experiences, to the effect that they are as one experiences them as being, 'by demonstrating' the relevant objects and properties when one experiences them, as Thau admits, then why cannot one introduce names for those objects and properties? If one could do this, one could close one's eyes and think a thought about just those objects to the effect that they have just these properties by using these names. Since the identity of the content of the thought, for the neo-Russellian, is determined simply by the objects thought about and the properties assigned to them (and their quasi-logical arrangement, which it there is no reason to think has to be any different), the thought one then thinks would have the same content as the thought one had when one's eyes were open. But this contradicts Sufficiency. Thau is thus driven explicitly to deny the possibility of introducing names for the properties one immediately experiences things as having. But this seems entirely unmotivated.²³

These problems would arise no less severely if one combined content conceptualism and Sufficiency with an even more coarse-grained view of content, such as a Lewis-Stalnaker type view. If one wants to combine content conceptualism and Sufficiency one ought, then, to settle for a more fine-grained view of representational content, such as some sort of neo-Fregean view. Within the framework of such a neo-Fregean view, one might try to make sense of Sufficiency in terms of the idea that the modes of presentation of objects and properties that compose the content of experience essentially have some sort of 'phenomenal glow' to them: they somehow bear or necessitate the phenomenal character of the experience in question. If this neo-Fregean theorist wants

²³ This criticism of Thau is well developed in Braun (2004).

to be a content conceptualist, he must grant that one can have a Wall-Thought with exactly the same content as Wall-Experience (as long as one is in a state of mind with the same phenomenal character as Wall-Experience). However, unlike Thau, he does not need to deny that if one closes one's eyes, one can have a thought about the wall that ascribes to it just the properties it was experienced to have (and indeed thought to have, in having Wall-Thought). The critical point for him is that this thought, entertained in the absence of Wall-Experience, cannot have just the same representational content as Wall-Thought, since it is prevented from involving the same 'phenomenally glowing' modes of presentation as Wall-Thought does.²⁴

While this combination of content conceptualism, Sufficiency, and neo-Fregeanism is somewhat more plausible than the combination with neo-Russellianism in place of neo-Fregeanism, it is still unattractive. In particular, the content conceptualism part of the package is unmotivated given the two other elements. Sufficiency, as we noted, was

²⁴ The notion of a 'phenomenally glowing' mode of presentation gestured at here is not to be equated with the notion of a 'phenomenal concept' or 'phenomenal mode of presentation' used in Essay 5 below. The notion of a phenomenal concept used in Essay 5 is the notion of a concept grasp of which requires that one has certain correlated species of phenomenal knowledge, viz. that one knows what it is like to have certain experiences or (as will be argued in Essay 5) what certain perceptually apparent properties are like. However, it is *not* built into the notion of a phenomenal concept that whenever one is in a state where these phenomenal concepts are used one actually is in a state with the phenomenal character associated with these concepts. For example, I can truly affirm to myself at a certain point in time that I know what it is like to experience beige, even if I do not at that time actually see anything beige or imagine beige. (Perhaps knowing what it is like to experience beige requires that *ability* to sensorily imagine beige, but one can have this ability without exercising it at a given time.) If the idea behind the notion of phenomenal concepts is right, my knowledge of what it is like to experience beige involves the use of a phenomenal concept. The notion of a 'phenomenally glowing' mode of presentation, by contrast, is the notion of a mode of presentation such that whenever one is in a state of mind involving it one must be in a state with a specific correlated phenomenal character. On the other hand, it is not built into the very notion of a 'phenomenally glowing' mode of presentation that being in a state of mind with such a mode of presentation requires having correlated phenomenal *knowledge*. Perhaps (supposing the notion of a 'glowing' mode of presentation is otherwise coherent) being in a state with such an mode of presentation will always go together with having certain correlated phenomenal knowledge, but this is a further claim, not built into the notion.

supposed to be explained in terms of the idea that the modes of presentation of objects and properties that compose the content of experience somehow bear or necessitate the phenomenal character of the experience. If this is how we conceive of the content of experience, then *prima facie* the plausible conclusion is that the content of experience is *not* the content of a possible belief or thought. If thoughts cannot have such ‘glowing’ contents, while experiences do, this would at least give some sort of account of the phenomenological contrast between experience and thought. If thoughts can have ‘glowing’ contents too, as content conceptualism would have us say, then why should there be this phenomenological contrast? This would seem to be an embarrassing question for the current package.

On both a neo-Russellian, a Lewis-Stalnakerian, and a neo-Fregean view of content, denying Insufficiency is an unattractive option for the content conceptualist. Since these are the major views of content, I conclude that the content conceptualist is not well advised to take it.

Denying the Thesis of Diaphanousness, (D)

The obvious remaining alternative for our conceptualist would be to deny (D), the thesis we have singled out as *the* thesis of diaphanousness (somewhat arbitrarily from the nexus of considerations of diaphanousness with which it is associated). Now, my aim in this paper is not to defend (D); such a task would take us far beyond its scope. Indeed, I am prepared to allow that (D) may ultimately stand in need of correction. What would be a greater worry is if (D) proved obviously stronger than, or less plausible than, the various claims of diaphanousness that have been so popular in recent philosophy of consciousness. If this were so, the conclusion that (D) is incompatible with content conceptualism would appear much less interesting than one might at first have thought.

This does not seem to be so, however. As I remarked above, Harman (1990) certainly seems to be committed to (D) in his influential argument against the idea that we need to posit any intrinsic qualitative properties of experiences.²⁵ In another recent, influential discussion of experience and diaphanousness/transparency, Michael Tye writes:

When we introspect our experiences and feelings, we become aware of what it is like for us to undergo them. But we are not directly aware of those experiences and feelings; nor are we directly aware of any of their qualities. The qualities to which we have direct access are the external ones, the qualities that, if they are qualities of anything, are qualities of external things. By being aware of these qualities, we are aware of phenomenal character. (2000: 51)²⁶

In saying that we are aware, in introspection, of the phenomenal character of an experience ‘by being aware of these qualities’ (viz. the qualities external things are represented as having) Tye is reasonably read as saying that it is *only* by being aware of these qualities that we are aware of phenomenal character. Otherwise Tye’s remark is much less interesting than it has standardly been supposed to be, as it then leaves open that we may be aware of the phenomenal character of experience in plenty of other indirect ways too. If it left this open, the claim could not take the significant argumentative burden Tye puts on it. Now, since the qualities external objects are

²⁵ See Stoljar (2004) for detailed discussions of Harman. Stoljar interprets Harman as committed to the claim that ‘[i]n introspection, one is or becomes aware of the intrinsic features of one’s experience *by* attending to the objects and properties represented by that experience.’ As I discuss in a moment, someone who is committed to this thesis (which I will call (S)) ought to accept (D).

²⁶ I am indebted to Stoljar (2004) for drawing my attention to this passage. It is clear from context that ‘these qualities’ in the last sentence are the qualities external objects are represented as having in the relevant experience.

represented to have, for a representationalist such as Tye, is just those they appear to have, this means that Tye is committed to (D).

Another recent, more critical discussion of diaphanousness is that in Stoljar (2004). Although Stoljar is critical of several claims about and arguments from diaphanousness, he does not find a reason to dismiss the following:

(S) In introspection, one is or becomes aware of the intrinsic features of one's experience *by* attending to the objects and properties represented by that experience. (2004:10; emphasis in original. I have changed Stoljar's numbering)

Stoljar is at this point operating against a background where it is assumed, for the sake of argument, that the 'intrinsic features' of one's experience, talked of in (S), bear or constitute its phenomenal character. Someone who accepts it under this assumption will thus also accept that, in introspection, one is or becomes aware of the phenomenal character of one's experience *by* attending to the objects and properties represented by that experience. And again, it is reasonable to read (S) as implying that, in introspection, one is or becomes aware of the intrinsic feature of one's experience *only* by attending to the objects and properties represented by that experience. That is to say, it is reasonable to read (S) as purporting to make a restrictive claim, as one that purports to rule out that, in introspection, one is or becomes aware of the intrinsic features of one's experience by doing F, where doing F is something entirely different from attending to the objects and properties represented by one's experience. Thus someone prepared to accept (S), ought to be prepared to accept (D).

The fact that some recent, influential discussions of the considerations of diaphanousness either are committed to (D), or fail to find reason to reject (D) as opposed to some other claims of diaphanousness that they do find reason to reject, of course does not show that (D) is true. What it does suggest is that our conclusion that (D) is incompatible with content conceptualism is not uninteresting or trivial. Consequently, the task this conclusion leaves us with, of finding out whether this incompatibility should be blamed on (D) or on content conceptualism, is not uninteresting or trivial.

Object-Directed Knowledge Arguments and Phenomenal Concepts in Perception

0. Introduction

Frank Jackson's (1982, 1986) knowledge argument aims to refute the physicalist thesis that all facts about the world are physical facts. According to a popular, and, I believe, plausible, view of the argument, one of its most important lessons is nevertheless that certain special introspectively acquired types of knowledge of experience involve certain special concepts or 'modes of presentation' of properties of experiences. These special concepts or modes of presentation deserve to be called 'phenomenal' notably because grasping them presumes having correlated phenomenal knowledge – i.e. presumes knowing what it is like to have experiences with the properties these concepts refer to. Many of those who take this 'perspectivalist' view of the argument maintain that these properties themselves are physical, and thus that Jackson's argument ultimately is unsound. However, for our purposes here, physicalism is besides the point. For our purposes, we may rely instead on the weaker claim that all facts about the world are capturable in 'book learning', to use Lewis's (1986) phrase.

In this paper I consider a parallel to Jackson's argument, differing from it in being, in a sense to be explained, an 'object-directed' rather than an 'experience-directed' knowledge argument. I argue that if the perspectivalist view of Jackson's experience-

directed argument is correct, then a parallel perspectivalist view of this object-directed knowledge argument is correct. This parallel view has the consequence that basic perceptual judgements about the world – say, judgements about what tomatoes are like, colour-wise, immediately based on visual experience – involve certain distinctive concepts. These distinctive concepts again deserve to be called ‘phenomenal’, since grasping them presumes having certain essentially experiential knowledge of what things having the properties these concepts refer to are like. We may think of these as ‘perceptual-phenomenal’ concepts, by analogy to the ‘introspective-phenomenal’ concepts that the perspectivalist view of the experience-directed knowledge argument is committed to.

The structure of the paper is as follows. In the next section, I briefly rehearse Jackson’s knowledge argument, and distinguish two popular types of responses to it: the ability and acquaintance hypotheses and the perspectivalist account. In section 2, I set out in some detail an argument showing how a commitment to phenomenal concepts in introspective judgement arises from giving a perspectivalist account of Jackson’s argument. In section 3, I briefly outline an object-directed counterpart to Jackson’s experience-directed knowledge argument. I then develop an argument showing how a commitment to phenomenal concepts in perceptual judgement arises from giving a perspectivalist account of the object-directed knowledge argument. This argument very closely parallels that set out in section 2. In sections 4 and 5, I reply to objections to two key steps in the argument in the previous section. I conclude, in section 6, that (conditionally on a perspectivalist view of the experience-directed knowledge argument being sound) basic perceptual judgements involve phenomenal concepts. I also add a brief comment on the possible significance of this conclusion.

1. Jackson's knowledge argument and the perspectivalist response

Jackson's knowledge argument aims, as we noted, to refute the physicalist thesis that all facts about the world are physical facts.¹ The argument runs as follows. Mary, a brilliant scientist, has since birth lived in an environment where everything is black or white, including her own body. She has never had any chromatic experience. Via black-and-white books, television, etc., Mary has received an extremely thorough education in physics, chemistry, biology, etc. She has learnt no less than the complete and perfectly accurate physical theory about colour perception, and indeed about everything else. Since each physical fact about the world is stated by this theory, there is no physical fact about the world Mary doesn't know. However, Jackson argued, there is something Mary doesn't know. Only when she is allowed to see something red will she come to know what it is like to see red. To come to know what it is like to see red is to come to know a further fact about experiences, a fact not included among the ones Mary already knew. This further fact about experiences must, then, be a non-physical fact. So the physicalist thesis is false

A great number of responses have been offered to Jackson's knowledge argument over the last twenty-odd years. Two popular types of response concede that Mary gains new knowledge when she first sees red, but that her new knowledge of what it is like to see red is not new propositional knowledge, but rather, respectively, new know-how or new acquaintance. These claims are known, respectively, as the 'ability hypothesis' and the

¹ Strictly, as Jackson (1982, 1986) mostly formulates matters, the physicalist thesis his argument purports to refute is that all information about the world is physical information (though, in (1986), he occasionally says that the argument shows there to be 'non-physical facts'). However, in the ensuing discussion of his argument, it has very often been reformulated as one aiming to show that not all facts about the world are physical facts. I follow this fairly standard reformulation here.

‘acquaintance hypothesis’.² Now, hardly anyone denies that Mary acquires new abilities (to imagine and to recognise experiences, perhaps among other things) and new acquaintance (with certain phenomenal features) on first seeing red. What is more controversial is whether her new abilities or new acquaintances can account for all of the new knowledge she gets when she first sees red. On the face of it, Mary does learn something of propositional shape when she first experiences red. After all, it is natural to think that when Mary first gets the chance to look at a tomato under normal chromatic conditions of perception, she may say to herself ‘So *this* is what it is like to experience a tomato (directly in daylight).’ To utter this sentence would seem a natural way for Mary to express what she learns in this case. The sentence thus uttered on the face of it expresses a proposition. The sentence in any case embeds in a normal way under propositional operators such as conditionals. Mary might go on to reflect to herself, ‘If *this* is what it is like to look at the sunset, it is makes sense that so many people are so keen on it.’ This suggests that a non-propositional construal of Mary’s utterance will come up against many of the problems that notoriously afflict non-cognitivist interpretations of moral utterances.³ At the very least, it seems that the default options should be that Mary’s utterance of the sentence ‘So this is what it is like to experience red’ expresses a proposition. If the utterance also, as we said, expresses what Mary learns when she first sees red, it is plausible to conclude that what Mary learns consists, in part, of new propositional knowledge.

Another popular response to the knowledge argument aims to respect this claim, while resisting or at least leaving open the conclusion that Mary comes to know a new, non-

² See Nemirow (1980) and Lewis (1986) for two early defences of the ability hypothesis. See Bigelow & Pargetter (1990) for a defence of the acquaintance hypothesis.

³ Loar (1990/1997) furnishes this objection to the ability hypothesis.

physical fact. The key idea of this account is that one and the same fact can be known in different ways – because it can be thought about in different ways. The fact can be grasped in thinking under different concepts, or via different modes of presentation, to use Frege's (1892) term. This account goes together with a coarse-grained conception of facts, on which the fact that Phosphorus is visible in the morning just is the fact that Hesperus is visible in the morning, and the fact that the glass contains H₂O just is the fact that it contains pure water. Facts, on this coarse-grained conception, are in effect true neo-Russellian propositions.⁴ States of belief, and states of propositional knowledge, are not, however, individuated simply in terms of the attitude of believing, or knowing, on the one hand, and a neo-Russellian proposition, on the other hand. They are also individuated in terms of concepts or modes of presentation of the objects and properties that figure in the neo-Russellian proposition.⁵ This is why it is possible to know that the glass contains pure water without knowing that it contains H₂O, or vice versa, despite the same fact being known to obtain in each case.

According to the present account, the correct statement of physicalism about facts is as the thesis that all facts *in this coarse-grained sense* are physical facts. So stated, the physicalist thesis is consistent with claim that Mary acquires new propositional knowledge when she has her first experience of red. Just as we cannot infer that water is not H₂O from the fact that one can know the glass contains water without knowing that it contains H₂O, we can't infer that Mary comes to know a new non-physical fact from

⁴ See Salmon (1986) and King (2005) for more on neo-Russellian propositions.

⁵ In what follows, for brevity's sake I shall write only of 'concepts', rather than 'concepts or modes of presentation'. I remain neutral on the exact relationship between concepts and modes of presentation, as I am on what either of these things are, except for the constraint (i) that they are things that allow for a finer individuation of beliefs than simply in terms of the attitude of believing and the neo-Russellian proposition that encodes the truth-conditions of the belief, and (ii) that they have a certain theoretical role in accounting for 'Frege cases' (cases of informative identity).

the premise that she comes to have new propositional knowledge about her experiences. If physicalism is correct, she only comes to know about an old fact about experiences in a new way. This account of what goes on in Jackson's knowledge argument is generally referred to as the 'perspectivalist' or 'phenomenal concepts' account.⁶ In my view, it remains the most promising response to the argument for physicalists.⁷

A salient feature of the perspectivalist or phenomenal concepts account is, unsurprisingly, its commitment to certain special concepts: phenomenal concepts. We might distinguish two sorts of reasons for calling these special concepts 'phenomenal'. The first, less important, reason has to do with the reference of the relevant concepts, with what properties these concepts are concepts of. The idea is that these concepts are concepts of the phenomenal properties of experiences, i.e. the properties or experiences that are or constitute their phenomenal character. If physicalism is right, these phenomenal properties are a species of physical properties. In what follows, I set this reason to count something a phenomenal concept (i.e. the fact that it refers to a phenomenal property of experiences) aside, and will consider it neither necessary nor sufficient to being a phenomenal concept. The second, more important, reason has to do with a necessary condition for possessing or grasping the concepts in question, namely that one has to have certain correlated phenomenal knowledge, to know what it is like to have certain experiences. The experiences in question are of course just those with the phenomenal properties these concepts are concepts of. It is this second feature of phenomenal concepts we will be concerned with in this paper.

⁶ The account is defended by Loar (1990/1977) and Lycan (1995, 2003) among many others.

⁷ I would like to stress, however, that the perspectivalist response to the knowledge argument does not presume physicalism, although most of the theorists who have developed and defended versions of it have been physicalists. Nor is physicalism an assumption for any of the arguments to follow. For the purposes of these arguments, the weaker assumption that all facts about the world are capturable in book learning suffices, as I discuss below.

In order to have a clear comparison for the object-directed argument we will turn to below, it will be useful to set out explicitly how the commitment to phenomenal concepts in this sense arises from the perspectivalist account of Jackson's knowledge argument.

2. An experience-directed argument for phenomenal concepts in introspection

I will now set out step by step how the commitment to phenomenal concepts arises from the perspectivalist account of Jackson's argument. Each step will be followed by comments (some brief, some longer) on its motivation.

- (1) Before seeing anything red, Mary knows all the (coarse-grained) facts about experiences of red, and indeed everything else.

One way to motivate (1) is by appeal to (i) the physicalist thesis that all facts about the world are physical facts, in conjunction with (ii) the claim that the physical facts about the world are just the facts that are stated by a complete and perfectly accurate theory of the world, (iii) the assumption that such a theory can, in principle, be imparted to a subject simply by black-and-white means, and (iv) the idea that a subject who otherwise is pretty much like us, viz. Mary, could in principle grasp such a theory. It is worth noting, however, that reasons for (1) can also be given independently of the physicalist thesis. An alternative motivation for (1) is the idea that all the facts about the world can be encoded in 'book learning', to use Lewis's (1986) phrase, i.e. can be stated by a complete and perfectly accurate theory that could be written down in a book, or, anyway, imparted to a subject by black-and-white means. If there are non-physical facts about the world, but all of these non-physical facts are capturable in book learning, the case for (1) remains no less strong. Physicalism is thus a stronger thesis than we strictly

speaking need to motivate (1). This is important since it means that also those dualists who believe that all the non-physical facts are capturable in book learning (and who reject the ability/acquaintance-hypotheses) are committed to phenomenal concepts in introspective judgement. If, in the following, I speak as if I am assuming physicalism, a proviso about the possibility of a weaker justification for (1) should be taken as tacit.

- (2) Thus, before seeing anything red, for any fact f about the world, Mary believes f to hold.

This step is justified by the idea that knowledge entails belief.

- (3) Before seeing anything red, Mary is not in a position to know what it is like to experience red.

This is a basic intuition underlying Jackson's knowledge argument.

- (4) After seeing something red, Mary does know what it is like to experience red.

This is another basic intuition underlying the argument.

- (5) Mary's learning what it is like to experience red involves her acquiring a belief, E , which she could appropriately express with the words 'So *this* is what it is like to experience tomatoes (directly in daylight)'.

This step is plausible once we give up the ability and acquaintance hypotheses, and settle for what we might call 'phenomenal cognitivism'.

- (6) (i) There is a neo-Russellian proposition p that constitutes the structured truth-condition of E ; moreover (ii), Mary had a belief M with the same structured truth-condition p before seeing anything red.

This step makes two claims. Claim (i) rests on the assumption that any belief has a structured truth-condition that takes the form of (or anyway can be represented by) a neo-Russellian proposition. I believe this assumption is plausible. It also greatly

facilitates the formulation of the arguments to follow. It is, however, strictly inessential to them, but reformulating them without it complicates matters considerably.

The case for (ii) is this. Assume, as is plausible, that E is true. Then p is true. Since facts are true neo-Russellian propositions, p is among the facts about the world. It follows from (2), then, that Mary believes p to hold. To believe p to hold is to have a belief whose structured truth-condition p is (or represents). So Mary had a belief M with the same structured truth-conditions as E.

Assume, on the other hand, that E is not true. It follows that Mary did not get any new propositional knowledge about what it is like to experience red, for E has the same content as this supposed propositional knowledge (E would be entailed by that knowledge if she had it). But it is part and parcel of the perspectivalist response, as standardly understood, to construe (4) so as to entail that Mary acquires new propositional knowledge of what it is like to experience red. Quite apart from that point, to assume that E is not true is to commit oneself to an error-theory about phenomenal judgement. Error-theories are generally comparatively unattractive theoretical options, for semantical and other reasons. It seems particularly unattractive in this case, as it is hard to identify any illusory experience or other types of misleading evidence that could account for the error in question, since E *just is* a judgement about experience.⁸

Moreover, even if E is false, a case for (ii) can be made. Suppose that a conceptually and linguistically sophisticated subject other than Mary, with normal chromatic experience, has E. Let's call this subject Bill. E has the structured truth-condition p ,

⁸ Admittedly, one could, in principle, appeal to a misleading higher-order experience of the first-order experience to explain the misattribution of a certain property to the latter.

where p is a quasi-logical or quasi-linguistic structure of objects and properties. The plausible assumption is that Bill, when having E , is in a position to introduce proper names referring to the relevant objects and predicates referring to the relevant properties. Using these names and predicates, he can put together a suitably grammatically structured sentence S whose structured truth-condition is p . To fix ideas, let's suppose that S is simply the sentence 'Experiences or red are röööd.' As it happens, the term 'röööd' catches on in Bill's speech community (which is just a sub-community of the English-speaking world). Bill and his peers talk with each other about what experiences 'röööd' embraces, and what experiences it doesn't. We may even assume, if we want, that some subjects are generally regarded as more expert than others when it comes to judging whether 'röööd' applies to experiences, and that some of Bill's peers defer to them. The key point is that 'röööd' figures as a predicate in a public language. Now, Mary is very much a part of Bill's speech community; indeed, Mary and Bill chat relentlessly over the telephone, black-and-white telecom, etc. These conversations are only exceptionally afflicted by any apparent failure of understanding on the part of either Mary or Bill. From time to time, Bill remarks that experiences of red are röööd. Mary picks up on this use of the word röööd, and begins using it as a predicate in her public language. Perhaps she does not count as an expert on the application of the term. But, if she does not, it does not mean that she cannot have attitudes in the specification of the contents of which the term 'röööd' may properly be used. Similarly, Bill is no expert on the application of such terms as 'arthritis' or 'beech', yet, as Burge (1979) argued, this does not mean that he cannot have attitudes in the specification of the contents of which those terms may properly be used. In fact, Bill believes on the strength of what he takes to be reliable testimony that there's a lot of beech around, and that old people often suffer from arthritis. Likewise, there is no obstacle to our

supposing that Mary believes, on the strength of what she takes to be reliable testimony from Bill, that experiences of red are röööd. Since the sentence ‘Experiences of red are röööd’ by assumption has the structured truth-condition p , Mary’s belief that experiences of red are röööd is a belief M whose structured truth-condition is p . Thus, even if E is false, a case for (ii) can be made.

Continuing the argument, the next step is

- (7) E is different from M .

This is a simple consequence of Leibniz’s Law, given that E is a belief Mary acquires only after seeing red, while M is a belief she had before.

- (8) For any beliefs x and y , x is different from y , if and only if either x and y differ in the neo-Russellian proposition that gives the structured truth-conditions of the beliefs, or in the concepts the beliefs involve.

This seems to be a plausible condition of individuation for belief states (considered as universals), at least as long as we permit ourselves to abstract away from differences stemming from different degrees of belief.⁹

- (9) E and M differ in the neo-Russellian proposition that gives the structured truth-conditions of the beliefs, or in the concepts the beliefs involve.

This obviously follows from (7) and (8).

⁹ Notice that (subject to a minor qualification) (8) is acceptable both to neo-Fregeans who take concepts (understood as modes of presentation) to be constituents of the propositional content of belief, and to those neo-Russellians who hold that, even though concepts are not constituents of the propositional content of belief (except in the special cases where the belief is a belief about concepts), belief is underlain by a mediated, belief-like relation to propositional content *by means of* concepts. Neo-Russellians taking the latter view include Salmon (1986), Soames (1987) and Braun (1998). (The minor qualification is that not all of these neo-Russellians will call the relevant, ternary belief-like relation ‘belief’. Salmon, for example, introduces the technical term ‘Bel’ for this relation, and takes belief to be the existential closure of Bel. (8), and the arguments in which it is used, could be reformulated so as to take this minor qualification into account, but only at the cost of considerable complications, so I rest content with pointing out their possibility.)

(10) E and M differ in the concepts the belief involve.

This obviously follows from (9) and (6)

(11) Mary was in no position to have E (unlike M) before knowing what it was like to experience red.

E is a belief that is entailed by, and has the same content as, the propositional knowledge that *ex hypothesi* constitutes knowing what it is like to experience red. Some if someone knowledgeably has E, they know what it is like to experience red. Now, if Mary was in a position to have E in her black-and-white room, she was in a position to have that belief communicated to her by testimony, to acquire it through black-and-white television or ditto sensory experience, or arrive at it by a priori reasoning. But if she is in a position to arrive at E by either of these means, she is in a position to knowledgeably have E. There is no reason to suppose that she can get E through testimony, or black-and-white sensory experience, or a priori reasoning *only when* these methods of belief acquisition operate unreliably. So if she has in a position to have E in the black-and-white room, she was in a position knowledgeably to have E in the black-and-white room. But this is to say that she was in a position to know what it is like to experience red before she was in a position to know what it is like to experience red, since she, by assumption, was not in any such position in the black-and-white room. So (11) is true.

(12) E (unlike M) involves a concept a subject is in no position to grasp unless she knows what it is like to experience red.

This claim explains the distinctive feature of E, as compared with M, affirmed in step (11) in terms of the difference between E and M affirmed in step (10). Since E and M do not differ in their neo-Russellian propositional content, and since beliefs with the same neo-Russellian propositional content are the same except insofar as they differ in

the concepts they involve, as per (8), the difference in concepts between E and M is the only explanation of the distinctive feature of E, as compared with M, affirmed in (11). So (12) is true. Given what we distinguished as the reason to count something a phenomenal concept, it follows from (12) that E involves a phenomenal concept.

Now, E is an introspective belief or judgement. This is not because its subject matter concerns the properties of experiences. After all, Mary has plenty of beliefs about the properties of experiences of red drawn from her neuroscience tutorials that are not introspective. E is introspective since it is arrived at through introspection.¹⁰ In the next section, we shall consider a parallel argument for positing phenomenal concepts in normal, ‘extrospective’ perceptual judgement.

3. An object-directed argument for phenomenal concepts in perception

Jackson’s knowledge argument is what we may call an ‘experience-directed’ knowledge argument. That is to say, the new knowledge Mary is alleged to acquire, when she comes to have experiences of the right sort, is explicitly about experiences. It is knowledge of what it is like to *experience* red. However, variants of the knowledge argument can be formulated on which the new knowledge Mary has purportedly acquired, when she comes to have experiences of the right sort, is not explicitly about

¹⁰ Note that I use ‘introspection’ in a relatively neutral way here. To say that E is arrived at by introspection is to say that it is reached by that apparently non-inferential and privileged way each of us has of finding out about our own mental states. (It is privileged not in that it is infallible, nor in that it necessarily has a comparatively high epistemic standing (compared with normal extrospective perceptual judgements, say), but in that others cannot find out about our mental states in the same way.) It is not part of the meaning of the very term ‘introspection’, on the neutral use, that it involves some form of inner perception (though the etymology of the term of course suggests this). As we shall see, the ‘inner sense’ theory is but one theory of introspective access.

experiences, but rather about the *objects* of those experiences. We may call arguments of the latter type ‘object-directed’ knowledge arguments.¹¹

In broad outline, the argument may run as follows. The initial assumptions about Mary remain the same: she has spent all her life in the black-and-white environment, and become physically omniscient. One day she is shown a ripe tomato under normal chromatic conditions of perception. Upon seeing it, Mary learns something. She comes to know what tomatoes are like, with respect to a certain visually salient aspect of their surface. She comes to know what tomatoes are like, colour-wise. To come to know this is to come to know a new fact about tomatoes, a fact not included among the physical facts about tomatoes. Thus, perhaps surprisingly, physicalism about tomatoes is false.

If we prefer a perspectivalist response to the experience-directed knowledge argument, we ought, I will argue, to prefer a perspectivalist response to an object-directed knowledge argument of this sort. That is to say, if we take the perspectivalist line in the former case, we should allow that Mary gets new propositional knowledge about tomatoes, but deny that the new propositional knowledge entails that a new fact is known. If physicalism is correct, Mary only comes to know of an old fact about tomatoes in a new way. Therefore, if physicalism is correct, Mary’s new perceptual knowledge of the tomato involves distinctive concepts of the tomato’s surface properties. These concepts are distinctive in that Mary couldn’t have grasped them before having some essentially experiential knowledge of what the certain things are like. Mary’s perceptual judgement, in other words, involves a type of phenomenal

¹¹ Variants of object-directed knowledge argument closely related to that to be given in the text are discussed by Byrne (2006) and Thau (2002). (The arguments of this paper were written and presented before Byrne’s paper came to my attention.) Campbell (2002: ch. 12) discusses another type of object-directed knowledge argument, focusing on our knowledge of categorical properties.

concepts. Or so we ought to agree if we agree with the perspectivalist account of the experience-directed knowledge argument.

I shall now attempt to make the case for this conclusion in finer detail. I will set out step by step how the conclusion arises from giving a perspectivalist response to an object-directed knowledge argument. The steps will correspond closely to the steps of the experience-directed argument in section 2. As there, we append comments (some brief, some longer) to each step on their motivation.

- (1) Before seeing the tomato, Mary knows all the (coarse-grained) facts about tomatoes, and indeed everything else.¹²

This again is motivated by physicalism about facts (or, alternatively, by the weaker claim that all facts about the world can be stated in book learning).

- (2) Before seeing the tomato, for any coarse-grained fact *f* about tomatoes, Mary believes *f* to hold.

Because knowledge entails belief.

- (3) Before seeing the tomato, Mary is in no position to know what tomatoes are like, colour-wise.

This step is likely to be considerably more controversial than its counterpart in the experience-directed argument. It is natural to object that (3) is false since Mary is in a position to know that tomatoes are red, and knowing this is to know what tomatoes are

¹² For present purposes, 'seeing the tomato' is short for 'seeing the tomato under normal chromatic conditions of perception'. Seeing the tomato on black-and-white television, in very poor lighting, etc. does not count as seeing it. Remember also that the experience of the tomato is Mary's first chromatic experience. There is no suggestion that she can imaginatively extrapolate to it from having seen, say orange and purple (if such a thing is possible).

like, colour-wise. We shall examine this objection in the next section; for now, let (3) stand as a supposition.

(4) After seeing the tomato, Mary knows what tomatoes are like, colour-wise.

We assume, of course, that Mary's chromatic colour vision operates normally the first time she sees one, and that all other conditions of chromatic perception are normal, or, if one wishes, optimal. Mary is therefore in a really good position to know what they are like colour-wise.

(5) Mary's learning what tomatoes are like, colour-wise, involves her acquiring a belief, O, which she could appropriately express with the words 'So tomatoes are like *that*.'

Part of the motivation for this claim is, again, the rejection of the ability or acquaintance hypothesis about what Mary learns when she first sees the tomato, given the truth of steps (3) and (4). Part of the motivation, that is to say, is the assumption that Mary's learning involves her acquiring propositional knowledge or at least a new belief. Even if one accepts this, however, one might be sceptical of the suggested identification of the content of the new belief Mary acquires, in terms of the words 'So tomatoes are like *that*'. Someone might object that even if one accepts what we called 'phenomenal cognitivism' above, the new belief O she acquires is not interestingly different from the belief E she was found to acquire in the experience-directed argument. Since different forms of words are given as expressive of O in (5) here than were given as expressive of E in the counterpart of (5) in the experience-directed argument, this means that (5) is false or at best misleading. Again, this objection is one I shall return to in the next section. For now, (5) stands as an assumption.

- (6) (i) There is a neo-Russellian proposition q that constitutes the structured truth-condition of O ; moreover (ii), Mary had a belief R with the same structured truth-condition q before seeing the tomato.

The motivation for (i) here is as for the corresponding claim in the experience-directed argument. The motivation for (ii) is also parallel. If O is true, then q is true; so q is among the facts; so Mary believed q to hold before seeing the tomato; so she had a belief R with the same structured truth-condition q as O . If O is not true, then consider a linguistically and conceptually sophisticated subject, say Bill, who has had all the relevant experience and who now has O . Bill, it seems, can introduce names for the objects, and predicates for the properties, that figure in q , and so formulate a sentence whose structured truth-condition is q .¹³ (Bill may not need to introduce new names and predicates, if they already exist in the language. Whether or not they do, does not matter, as long as the terms one introduces are of a sort that allows them to catch on as common currency in the public language.) To fix ideas, let's suppose that 'Tomatoes are red' is such a sentence. The various terms in this sentence catch on as terms of the public language, of which Mary herself counts as a competent member. Speaking with Bill she comes round to the belief that tomatoes are red. Even if she does not count as, or is not treated as, an authority on the application of 'red', this does not mean she cannot have attitudes in the specification of the contents of which 'red' properly may be used, as Burge (1979) showed. Her belief that tomatoes are red is, *ex hypothesi*, a belief

¹³ Thau (2002) in effect denies precisely this step. He argues that the property a normal subject most directly sees a tomato to have, when he sees what the tomato is like colour-wise, is not a property anyone could give a non-demonstrative name for, no matter how linguistically competent they may be. See Essay 5 above, and Braun (2004) for critical discussion of this view of Thau's.

R with the same structured truth-condition as O. So whether O is true or not, a plausible case for (ii) can be made.¹⁴

The following steps are again parallel to the experience-directed case, and have entirely parallel motivation:

- (7) O is distinct from R. [LL]
- (8) For any beliefs x and y , x is different from y , if and only if either x and y differ in the neo-Russellian proposition that gives the structured truth-conditions of the beliefs, or in the concepts the beliefs involve. [Crit of individ. for belief]
- (9) O and R differ in the neo-Russellian proposition that gives the structured truth-conditions of the beliefs, or in the concepts the beliefs involve. [(7), (8)]
- (10) O and R differ in the concepts the beliefs involve. [(6), (9)]
- (11) Mary was in no position to have O (unlike R) before knowing what tomatoes are like, colour-wise. [As for (11) in the last section.]
- (12) O (unlike R) involves a concept a subject is in no position to grasp unless she knows what tomatoes are like, colour-wise. [As for (12) in the last section]

We thus get the conclusion that O involves a phenomenal concept, in the sense distinguished in section 1.

Now, whether or not the object-directed argument shows anything interestingly different from the experience-directed argument depends on the relationship between beliefs O and E. On the face of it, they differ in two ways. First, in subject matter, since

¹⁴ One difference between the case for (ii) here and the counterpart case for (ii) in the experience-directed argument is that error-theory, sometimes referred to as projectivism, has been considerably more popular. See Boghossian & Velleman (1989, 1991), Thau (2002) and Chalmers (2006) for error theories about various types of colour-properties. Since the case for (ii) goes through also on the assumption of error theory, this difference does not ultimately matter.

E is a belief about what *experiences* of red are like while O is a belief about what *red things* are like, in a certain respect. This difference in subject matter is reflected in what we have identified as the natural or canonical ways of expressing the beliefs: E in terms of ‘So *this* is what it is like to have a visual experience of tomatoes (directly in daylight)’; O in terms of ‘So tomatoes are like *this*, colour-wise’.¹⁵ Second, the beliefs differ in epistemic role. E is an introspective belief, arrived at by introspective reflection on one’s experience. O is a normal, ‘extrospective’ perceptual belief, arrived at in the normal perceptual way, on the immediate basis of perceptual experience where attention is directed upon the objects of experience.

If the steps of the argument above can be defended in a way that is consistent with upholding this difference between O and E, the object-directed argument has told us something ‘new’ and interesting, namely that phenomenal concepts are involved also in normal, ‘extrospective’ judgements about our surroundings. As we saw, there were in particular two steps in the argument that were likely to provoke objections. To these we now turn.

4. Does Mary already know what tomatoes are like, colour-wise?

A natural objection to (3) is that Mary is in a position to know that tomatoes are red while in the black-and-white room, and that knowing this much is knowing what tomatoes are red, colour-wise. One way of meeting this objection is to argue that Mary is not in a position to know that tomatoes are red prior to having first-hand experience of something red. As the broadly Burge (1979) inspired arguments above suggest, I am

¹⁵ The reason I formulate these utterances in terms of ‘tomatoes’ rather than in terms of ‘red (things)’ is to remain neutral on the question of whether Mary should be thought of as a competent user of the term ‘red’. As suggested above, I believe that we ought to think of her as such. Still, it is worthwhile to observe that we can find natural expressions of her beliefs here that do not presume that she is.

not terribly tempted by this reply. I am prepared to allow that Mary is in a position to know and believe tomatoes are red, while in the black-and-white room, through competent, deference-dependent use of the word 'red'. What I will argue instead is that this knowledge does not, in the intended sense, amount to knowledge of what tomatoes are like, colour-wise.

A necessary condition for having the latter knowledge is having is the ability visually-demonstratively to recognise redness. It requires that ability to judge, perceptually and non-inferentially, that red is *that* colour. Mary arguably lacks this ability. To make this vivid, suppose that Mary's supervisors decide to present her with the colour counterpart of Molyneux's Question. Mary is told a blue sphere and a red sphere will be presented to her. Since, as we are currently assuming, Mary knows some things are red and other blue, she will be able to understand of what she is thus told.¹⁶ The question is: when Mary is presented with the spheres, will she be able to tell which sphere is blue and which red?

An argument for a 'No' answer here is as follows. Suppose Mary doesn't know what red things look like, or what blue things look like. Then, surely, she cannot tell by looking which is the blue sphere and which is the red. Compare: if you don't know what Tony Blair looks like nor what George Bush looks like, and you see them together up at Gleneagles, you can't tell by looking who is who. Now it is plausible that Mary doesn't know what red things look like or blue things look like. A necessary condition for knowing what red things look like is knowing what way red things look. A necessary

¹⁶ Her understanding may well be incomplete, like my understanding of 'beech' surely is. But since incomplete understanding is commonplace, this does not mean that she in an ordinary sense understands what she is told.

condition for knowing what way red things look is knowing something of the form: red things look *F*, where '*F*' stands in for a suitable expression picking out (via a suitable concept) the way red things look. Thus a necessary condition for knowing what red things look like is knowing that they look *F*, given a suitable replacement for *F*. However, knowing that red things look *F* is *sufficient* for knowing what red things look like only when *F* is replaced by a demonstrative identification of the way in question.¹⁷ Thus Mary knows what red things look like iff she knows that they look *that* way (for a suitable demonstrative identification).

When Mary looks at the tomato, she surely knows that the tomato looks *that* way. If Mary knew what red things looked like, prior to seeing the tomato, the knowledge that the tomato looks *that* way, which she has when looking at the tomato, is not news to her. However, knowing that the tomato looks *that* way plausibly constitutes or anyway suffices for knowing what it is like to see something red, which *ex hypothesi* is news to her. On the other hand, if one can know that red things look *that* (demonstratively identified) way, but nevertheless fail to know what it is like to see something red, whether ordinary perceivers really know what it is like to see red becomes a dubious proposition, for it is dubious whether ordinary perceivers possess, or anyone possesses, what more would be required in order to know what it is like to see red. The supposition that Mary does not know what red things look like, prior to seeing anything red, is therefore highly plausible.

Still, someone may press the case for a 'Yes' to our counterpart of Molyneux's Question. The best case for answering 'Yes', it seems to me, appeals to the various

¹⁷ For example, if one knows that red things look red, but does not know that red things look *that* way, one does not know what red things look like.

asymmetries in colour space.¹⁸ Thus it is said that blue things tend to appear cool while red things tend to appear warm, for example. Mary may know that these generalisations hold. On the one hand, this knowledge does not imply that she knows what it is like to see red things, like the knowledge that red things look *that* way. On the other, it is not trivial, like the knowledge that red things look red. It allows her to reason as follows: the surface of this object appears cool while the surface of that object appears warm; hence this object is probably blue and that probably red.

However, precisely because it appeals to reasoning, this argument accepts that Mary doesn't know what red things are like, colour-wise, when she first sees one. Suppose we intervene at the stage of her reasoning where she concludes that the surface of this object appears cool. If we told her at that point that *that* colour was blue, it would be informative to her. If it wasn't, her reasoning would be redundant. By contrast, our remark is not informative to any perceiver visually familiar with blue (and having the concepts required to understand the remark). It was in order to mark this contrast that I required, for knowing what redness is like, the ability to judge *non-inferentially*, on the basis of perception, that that colour is red.

***5. Is the new belief O Mary acquires interestingly different from the
experience-directed, introspective belief E?***

Above, we identified an objection provoked by (5), centring on the sentence identified as a natural or canonical way of expressing the new belief O, viz. 'So tomatoes are like *this*, colour-wise'. The objection may be put as a dilemma. Either the correct interpretation of the sentence 'So tomatoes are like *this*, colour-wise' is one on which it

¹⁸ On these, see, e.g., Hardin (1988).

effectively means the same as ‘So this is what it is like to experience the colour of tomatoes’. If so, then O would just as naturally or canonically be expressed by the latter form of words. The difference between O and E would accordingly collapse. On this horn of the dilemma, (5) may well be true (indeed, if the objector gives a perspectivalist response to the experience-directed knowledge argument, he is committed to think that (5) is true). But so understood, the object-directed argument does not add anything interesting to the experience-directed argument.

The other horn of the dilemma is that the correct interpretation of ‘So tomatoes are like *this*, colour-wise’ is one on which it purports to say something more directly about the surface characteristics of tomatoes. If this is how the words expressing O are to be understood, however, it is just false that getting to have an chromatic experience of a tomato makes Mary acquire for the first time such a belief. If this objector is a perspectivalist, he admits that Mary gets the new propositional knowledge or belief attributed to her in the experience-directed argument. However, he contends, that is the *only* interesting new knowledge or belief Mary gets.¹⁹

I will call the perspectivalist theorist who puts forward this objection the ‘Merely Experience-Directed’ or MED theorist. In support of his claim, he may make the following point. When we justified (3) in the last section, we argued that Mary would not know what red things look like, before she had seen any. We also argued that

¹⁹ The qualification ‘interesting’ here restricts knowledge and beliefs to those whose truth posits a prima facie problem for physicalism. This excludes new knowledge or new beliefs acquired by Mary simply as a result of taking part in certain new events. For example, when she is presented with the tomato, Mary will know she is now standing face to face with a tomato. She couldn’t have had that first-person present-tense knowledge before, since she never stood face to face with a tomato before. But the truth-condition of that piece of knowledge is still unproblematically physical.

knowing what red things look like suffices for or constitutes knowing what experiences of red things are like. The MED theorist fully endorses these two claims. However, he infers from these two claims that (3), in the object-directed argument, is just a variant formulation of (3) in the experience-directed argument, viz.: of the claim that Mary doesn't know what it is like to see red.

However, I will argue that the MED theorist is wrong or unmotivated in claiming that the *only* interesting propositional knowledge and belief Mary acquires is experience-directed. First, the MED theorist accepts that Mary's knowledge of what it is like to experience red is matter of her knowing what red things look like. As we argued above, Mary knows what red things look like iff she knows that red things look *that* way, where '*that*' is an appropriate demonstrative identification. Thus Mary learns something she could express in the words 'The tomato looks *that* way'. This is what Chisholm (1957) called a 'non-comparative' looks report, and Jackson (1977) a 'phenomenal' looks report.²⁰ For any such report, of the form 'x looks F', it is appropriate to ask 'Is x F?' Mary knows this, and she also knows or at least believes that if conditions of perception are normal, and x looks F, in the non-comparative sense, then x is F. Mary may therefore reason as follows: 'The tomato looks *that* way. Conditions of perception are normal. So, probably, it is *that* way.' In this way she acquires belief and perhaps even knowledge that is not experience-directed, but has just the character (5) says it has, on the interpretation on which (5), according to the MED theorist, is false.

A second argument against the MED theorist is based on the nature of introspection. The MED theorist accepts that Mary gets new experience-directed knowledge,

²⁰ See essay 1 above for further discussion of non-comparative looks reports, and for defence of the claim that some basic looks reports count as such.

appropriately expressible as ‘This is what it is like to experience red’, when she sees the tomato. The new knowledge here is in a neutral sense introspective. I will now consider two major and representative theories of introspection, and argue that on each of these theories, the MED theorist’s view is either false or unmotivated.

The first type of theory maintains that introspective access to perceptual experiences essentially proceeds via a certain object-directed judgement, made on the basis of the perceptual experience. I will refer to theories of the former kind as ‘ascent-routine’ views. They receive canonical formulation in the following passage from Gareth Evans.

[A] subject can gain knowledge of his internal informational states in a very simple way: by re-using precisely those skills of conceptualization that he uses to make judgement about the world. ... He goes through exactly the same procedure as he would go through if he were trying to make a judgement about how it is at this place now, but excluding any knowledge he has *of an extraneous kind*. (That is, he seeks to determine what he would judge if he did not have such extraneous information.) [H]e may prefix this result with the operator ‘It seems to me as though...’. This is a way of producing in himself, and giving expression to, a cognitive state whose content is *systematically* dependent upon the content of the informational state, and the systematic dependence is a basis for him to claim knowledge of the informational state. (Evans 1982: 227-8)

Evans talks of access to internal informational states, but the view has been generalised to the case of access to the phenomenal character of experiences.²¹ On the basis of the ascent-routine view, as outlined by Evans, we can argue as follows. The MED theorist holds that Mary gets introspective knowledge of what it is like to see red when she has her experience of the tomato. He also holds that that’s the only interesting new

²¹ Byrne (2001) generalises Evans’s view in just this way.

knowledge she gets. According to the ascent-routine view, she could get that knowledge only by 're-using precisely those skills of conceptualization' she would use in making a judgement about the world, and then embedding the content of that judgment within an 'appears' context. The resulting, knowledgeable appearance judgement (in which the content of the object-directed judgement is embedded) constitutes, or anyway is necessary and sufficient for, her knowledge of what it is like to see red. The only plausible candidate for an object-directed judgement giving the embedded content of the appearance judgement seems to be one constituting new knowledge or at least new belief about those objects: it is one expressible as 'So tomatoes are like that (colour-wise)!' This is just as premise (5) says (on the interpretation on which, according to the MED theorist, it is false).

Perhaps the MED theorist would object that an old piece of knowledge, or an old belief, could do just as well in the role of the object-directed judgement that forms the basis of the ascent routine. Suppose Mary believed all along the tomato she is to be presented with is red. On seeing it, she judges it to be red, and progresses, in accordance with the ascent routine, to the judgement that it appears to her that the tomato is red. The latter constitutes – or anyway suffices for – her knowing what it is like to see something red. Or so the MED theorist might suggest.

The problem with this suggestion is that it can't explain why Mary didn't know what it was like to see red well before seeing anything red. By assumption, Mary knows and believes in the black-and-white era that the tomato is red. If she knows that, she also plausibly knows that the tomato appears red when seen under normal conditions of perception. For Mary knows that redness, like shape, is one of the manifest qualities of

(medium-sized) things, and that things appear coloured and shaped as they in fact are coloured and shaped under normal conditions. Thus Mary could suppose she was seeing a tomato under normal conditions of perception, and judge, within the context of that supposition, that the tomato appears red to her. Again, according to the ascent routine, that is sufficient for her to know, within the context of that supposition, what it is like to see something red, and there is no reason to think that that knowledge is destroyed when the supposition is lifted. The MED theorist's suggestion, then, implies that Mary knew what it is like to see red already in the black-and-white era, contrary to assumption.

At this point, the MED theorist may decide that the culprit is the ascent-routine view of introspection. He may deny that introspective access to the character of perceptual experience essentially depends on one's making an object-directed judgement on the basis of that experience. Rather, introspection provides as direct access to the character of mental states as vision provides to the manifest qualities of outer things. A view of introspection of this kind may fairly be described as an 'inner-sense' view.²² William Lycan, an influential defender of the view, puts the view as follows:

According to ['inner sense' or 'higher-order perception'] theories, a subject's awareness of her or his own mental state consists in a representing of that state itself. The 'inner sense' view has it that the representing is done quasi-perceptually, by a set of functionally specified internal attention mechanisms of some kind that scan or monitor 'first-order' mental/brain states. (Lycan 2003: 390)

²² The epithet 'inner-sense' is not meant to suggest that the view cannot acknowledge certain disanalogies between introspection and 'outer' perceptual awareness. It is often remarked, for example, that introspection does not have a phenomenology of its own. However, the view is committed at least to think that outer perception is a good model for introspection.

Now, a crucial point for our purposes here is that the inner-sense theorist takes outer sensory perception to be at least a model for the nature of introspection. This feature of the view, I argue, makes it an unattractive option for the MED theorist, as it makes his overall position unmotivated. Remember, the MED theorist endorses the perspectivist response to the original knowledge argument: he accepts physicalism, and he accepts that Mary (nevertheless) gets new propositional knowledge of what it is like to experience red when she first sees something red. He therefore has to explain how it is possible for introspection to teach Mary something new about the character of experiences despite Mary already knowing all the coarse-grained facts about them. He must assume that introspection somehow provides new knowledge about old facts, by providing the subject with new concepts – new ways of thinking – of these old facts. Specifically, he must assume that the quasi-perceptual relation of inner sense can and does provide the subject with new concepts of the (coarse-grained) phenomenal facts about his experience. Now, since inner sense is modelled on outer sense, the feature just ascribed to inner sense strongly suggests that outer sense perception – vision, audition, etc. – can and does provide the subject with distinctive concepts of the coarse-grained external facts about tomatoes or other objects of perception. Hence it strongly suggests that new forms of outer perception, such as those Mary enjoys when seeing the tomato, give her a new concept of certain surface properties of the tomato, and hence a new belief about the tomato. But this is precisely what the MED theorist denies. The bottom line is that the MED theorist who accepts the inner sense theory is committed to a wholly unmotivated asymmetry between inner and outer perception.

On each of the two views of introspection discussed, then, we have strong reason to reject the MED theorist's view. Perhaps there is another view of introspection that is

both independently plausible and more congenial to the MED theorist. However, since the two theories discussed represent the two major approaches to introspection, our discussion gives at least some inductive grounds for doubting that the MED theorist will find what he needs. We have reason, then, to treat the remaining obstacle to asserting premise (5) as disposed of.

6. Conclusion

I conclude then, that if the experience-directed argument (in section 2) for phenomenal concepts in introspective judgement is sound, then the object-directed argument for phenomenal concepts in perceptual judgement is sound. The experience-directed argument is sound if the perspectivalist view of Jackson's knowledge argument is correct. We have not argued at length for the correctness of the latter in this paper, apart from a few critical remarks about the ability and acquaintance hypotheses, on the one hand, and a few supportive remarks in favour of the idea that all facts about the world are capturable in book learning, on the other. However, I believe this perspectivalist account is plausible. Thus, I believe our object-directed argument for phenomenal concepts in perceptual judgement plausibly is sound.

If this is so, it raises several interesting further questions. In particular, there is currently a lively debate among adherents of the perspectivalist account as to whether perceptual or demonstrative concepts provide a good model for introspective-phenomenal concepts.²³ Some of these theorists argue that they do, others that they don't, and that we instead ought to adopt some sort of 'quotationalist' account of phenomenal concepts, on which they contain instances of the properties they denote, in a way that is analogous

²³ See the papers in part two of Alter & Walter (2007) for more on this, particularly the papers by Levine and Papineau.

the way in which the word ‘the word “Boston”’ contains the word it talks about, viz. ‘Boston’.²⁴ Prima facie, the finding that analogues of phenomenal concepts figure in extroceptive, perceptual judgement would appear to be pleasing to the former group of theorists. Even if this appearance is ultimately misleading, we may reasonably hope that assessing different views of phenomenal concept in light of the analogues we have found between perceptual-phenomenal and introspective-phenomenal concepts will be fruitful to the debate on phenomenal concepts and phenomenal knowledge.

²⁴ A quotationalist view is defended in Papineau (2002). Levine’s contribution to Alter & Walter (2007) defends a demonstrative view.

Are Only Mental States Intentional?

0. Introduction

Like many of his contemporaries, Franz Brentano (1874) was exercised by the question of whether the emerging science of psychology had any unified domain of inquiry. His great contribution to that debate was his revival of the notion of intentionality, which, he suggested, held the key to the unity of the domain of mental phenomena. Restricted to mental *states*, his proposal entails a claim I shall dub ‘Brentano’s thesis’:

Brentano’s thesis: All and only mental states are intentional states.

An influential objection to Brentano’s thesis is that intentionality isn’t necessary to mentality, since moods, or bodily sensations, or other states of mind regarded as ‘qualitative’, are not intentional. Over the last decade, a number of philosophers, including Michael Tye (1995, 2000) and Tim Crane (1998, 2001), have done much to rebut this objection, presenting several interesting and powerful considerations in favour of an intentional view of qualitative states. These arguments are central to Crane’s (2001) recent sustained defence of Brentano’s thesis.

While much ink has been lavished on the necessity of intentionality to mentality, little attention has been devoted to its sufficiency. Yet this question gives rise to a significant and underappreciated challenge to Brentano’s thesis. In this paper, I take as my point of

departure Crane's lucid and, I believe, fairly representative articulation of intentionality. I show that on a simpleminded but natural reading of some central passages in that articulation, the non-mental state of attracting a metal bar comes out as intentional. I then consider two types of response. One is to spell out intentionality in explicitly mental terms; this unacceptably trivialises one half of Brentano's thesis. Another is to ratchet up the necessary condition for intentionality; unfortunately, a condition strong enough to exclude all non-mental states is so strong that it becomes highly doubtful whether all mental states meet it. Or so I will argue.

1. A problem

Crane (ch. 1) finds two essential ingredients in the notion of intentionality: directedness upon an object and aspectual shape.¹ He elucidates them largely by describing how they are exemplified by certain paradigm intentional states, thoughts in particular. Directedness upon an object shows up here in that all thought is thought of something or other. For any thought, there is a (possibly non-unique) answer to the question 'what is thought about, in having this thought?' Whatever answers this question is the (or rather an) object upon which the thought is directed. Aspectual shape shows up in the fact that whenever an object is thought of, it is thought of in a certain way. If you are thinking of St. Petersburg, I am thinking of Leningrad, and a third person, listening to Shostakovich's Seventh Symphony, thinks of the place where this piece was composed, we are thinking of the same city, but most likely in different ways. We think of the city as an exemplar of different features. It figures in our thoughts under different aspects.

¹ In what follows, all references are to Crane (2001) except where otherwise noted.

However, certain simple non-mental states seem to share the features which were noted as paradigmatically exemplifying directedness and aspectual shape. For example, all attraction is attraction of something or other. For any state of attraction, there is a (possibly non-unique) answer to the question ‘what is attracted, in this state of attraction?’. Whatever answers this question is naturally regarded as an object of attraction. Moreover, whenever an object is attracted, it is attracted in a certain way: gravitationally, electrically, magnetically, or perhaps in some other way. If a metal ball, say, gravitationally attracts a metal bar, say, it attracts it as a thing having a certain mass. If the metal ball electrically attracts the metal bar, it attracts it as a thing with a certain charge. The ball attracts the bar here, in the different cases, as an exemplar of different features. It figures in the state of attraction under different aspects.

This shows that on a simpleminded but natural understanding of intentionality, viz. that provided by the features of thought described two paragraphs back, some non-mental are not excluded from the intentional.

2. A dismissive response

A somewhat dismissive reaction to this problem would be to argue that ‘directedness’ and ‘aspectual shape’ were only ever intended to describe what is presented to a certain *mind*, as such, and how it is presented to it. Some of Crane’s remarks are formulated in a way that might suggest that this is how he thinks of it:

Remember that, in general, the intentional object of a state S is what is given in an answer to the question, ‘what is your *mind* directed on when in S?’ (p. 80, emphasis added)

The thesis of aspectual shape is that, where states of *mind* are concerned, there is no such thing as a pure reference. All *mental* access to objects is ‘one-sided and dependent on a standpoint’... (p. 20, emphasis added. The quotation is from Frege 1892.)

So, one might argue, our iron ball’s attraction of the bar isn’t so much as a candidate for being an intentional state, as it isn’t a mental state.²

The current reaction articulates what it is for something to have intentionality in explicitly mental terms. This makes intentionality trivially sufficient to mentality, and one half of Brentano’s thesis quite a bit less interesting than one might have thought. Is this trivialisation an acceptable price to pay for the truth of the thesis? I think not. The point can be made in terms of a dilemma. The dilemma arises over the apparent parallels, made out in the last section, between such relational states as attracting an object and such paradigm intentional states as thinking of an object. Specifically, it arises over the following question: is it so that these apparent parallels are genuine but that, nevertheless, the relational state of attracting is not intentional *simply* for the reason that it is not a mental state?

Suppose we affirm this question. We might, then, introduce the term ‘schintonality’ for the features that account for the admitted genuine parallels between attractions etc. and intentional states. It is hard to see what objections there could be, on this view, against defining intentionality as having schintonality and being mental. But if intentionality is so defined, then in no interesting sense is ‘our conception of the mind ... unified by the idea of intentionality’, as Crane (p. 2) claims that it is. Even if Crane’s claim here does

² I do not mean to suggest (nor to deny) that Crane would endorse this reaction.

not strictly follow from Brentano's thesis, I believe Crane is right in suggesting that any defence of that thesis worthy of our attention ought to vindicate the claim that the notion of intentionality somehow unifies the domain of the mental. After all, the question of the unity of the mental domain was the underlying question that prompted Brentano to formulate his thesis in the first place.

Let's assume that this point is accepted. Perhaps someone would argue that it would, nevertheless, be interesting to find that intentionality (i.e. schintinality plus mentality) is *common* to *all* states of mind. This finding would not be uninteresting, but it would not in any meaningful sense unify the domain of mental states, since plenty of non-mental relational states also are schintonal. Analogously, it would not be uninteresting to find that all mental states are causally relevant. But it would not in any meaningful sense unify the domain of mental states, since (to put it mildly) plenty of non-mental states also are causally relevant.

Suppose we deny the question – suppose, that is to say, we affirm that either the apparent parallels are not genuine parallels, or the state of attracting is non-intentional but *not* simply for the reason that it isn't mental. If the state of attracting is non-intentional but *not* simply for the reason that it isn't mental, the current dismissive reaction has not told us the full story of *why* the state of attraction is not intentional. Something more needs to be said. If the apparent parallels we found are not genuine parallels, the dismissive reaction has not told us *why* the apparent parallels are somehow misleading. Something more needs to be said. In either case, then, the dismissive reaction is inadequate.

3. Failures of substitution and of existential presupposition

A more interesting reaction to the problem that a state of attracting a metal bar comes out as intentional invokes some of the more theoretical remarks Crane makes about directedness and aspectual shape, in particular, that it is characteristic of directedness that intentional states may be directed upon things that don't exist, and of aspectual shape that reports of intentional states are susceptible to substitution failures with respect to the terms that specify their objects.

The failure of existential generalisation is an expression of the fact that the objects of some intentional states do not exist; the failure of substitutivity is an expression of the aspectual shape of intentional states. (2001: 21)

The state of attracting a metal bar lacks both of these characteristics: our ball cannot attract a metal bar if there are none, and, if 'metal bar' has the same extension as 'elongated metal cylinder', then

(1) The ball attracts a metal bar.

is equivalent with

(2) The ball attracts an elongated metal cylinder.

One might hope, then, that taking these two characteristics to be necessary to directedness and aspectual shape, respectively, would exclude all non-mental states. Unfortunately, it won't. Consider the disposition to attract a pretzel-shaped piece of metal ('a metal pretzel' for short). Even if there are no metal pretzels, something may be disposed to attract a metal pretzel. And even if the extension of 'metal pretzel' is the same as the extension of 'passion for shrimp-flavoured ice-cream', i.e. the empty set, the true report

(3) The ball is disposed to attract a metal pretzel

is not equivalent to

(4) The ball is disposed to attract a passion for shrimp-flavoured ice-cream.

Another problem for the proposal is that, conjoined with Brentano's thesis, it commits one to say that there is no such mental state as non-epistemic seeing, in the sense of Dretske (1969). It commits one to saying that there is no such mental state as the basic visual states expressed by standard usages of reports of the form 'S sees F' where 'F' is a noun phrase that refers to or quantifies over concrete particulars, for instance:

(5) Jim sees a ghost

(6) Tom sees the president of Kazakhstan

(7) Tom sees Nursultan Nazarbayev.

If there are no ghosts, then on the standard usage of 'seeing a ghost', (5) cannot be true. On the same usage, since Nursultan Nazarbayev is the president of Kazakhstan, (6) and (7) entail each other. The commitment that there is no such mental state as non-epistemic seeing is costly, perhaps even fatal, to the proposal. Nonetheless, I will set it aside in what follows and concentrate on the problem posed by dispositions.

To exclude them, we need an even stronger condition necessary for intentionality. I will now consider such strengthened conditions for each of the two strands Crane distinguishes in the notion.

4. Directedness upon impossibilities

With respect to directedness, one suggestion might be that states exhibit directedness only if they are of a kind allowing them non-vacuously to be directed not only upon *non-existent* objects but also upon *impossible* objects. This does not exclude thoughts

and desires, since we may think of or desire not only something that does not exist, such as world peace, but also something that could not exist, such as wealth in dollars greater than the greatest prime. Dispositions do not, however, seem to meet this requirement. It might be argued that our ball may be disposed to attract, say, a round and square plate, since it might be argued that the vacuous truth of the following counterfactual is sufficient to ascribe the disposition: if a round and square plate were part of the same physical system as our ball, our ball would attract it. But just because that counterfactual is vacuously true, anything would vacuously have the disposition (if indeed the truth of the counterfactual suffices for having the disposition). Thoughts of or desires for the impossible are not vacuously had in this way. The non-vacuity condition in the necessary condition above thus ensures dispositions are excluded from the class of states with directedness, while thoughts and desires are not.

Yet it now becomes much more doubtful whether all mental states meet the condition. Consider pain experiences. Even if we allow that pain experiences can have a non-existent object, as in phantom pains, it is very hard to conceive of a pain the correct description of which would require us to take it to be directed upon an impossible object. Consider also simple visual or other perceptual experiences where attention is fixed, or not even engaged with respect to the things these experiences are experiences of. Here also it is far from clear whether there are any possible instances of such simple experiences that are directed upon impossible objects.³ I admit, of course, that this is not a foolproof justification for the claim that pain experiences, or simple perceptual

³ It is arguable that more complex, higher-level experiences may be directed upon impossible objects. For example, experiences of Penrose triangles, or those occurring in the Waterfall Illusion (see Crane 1988), may be describable in such terms. Yet these experiences seem to presume the engagement of attention.

experiences, fail the requirement, since the hardness of conceiving of such a pain or experience might reflect the poverty of our imagination rather than a limit on what is possible. But it seems to be overly bold to take on the commitment that pains and simple perceptual experiences have to meet the requirement.⁴ It is also telling that this necessary condition sets the standard for intentionality much higher than either Tye or Crane have even purported to attain in their arguments for the intentionality of bodily sensations and moods.

5. Varieties of substitution failure

With respect to aspectual shape, I will suggest not one but two strengthenings of the necessary condition for this property. They are strengthenings compared to what we implicitly assumed the property to require above. There we inferred from the non-equivalence of (3) and (4) that dispositions were not excluded from the class of states with aspectual shape. Now, the object-specifying terms in (3) and (4) have, ex hypothesi, exactly the same extension. However, this is a relatively weak semantic equivalence relation between two terms. Correspondingly, the type of substitution failure that the non-equivalence of (3) and (4) exemplifies is a relatively weak one. Perhaps it is *too* weak to be an instance of the kind of failure of substitutivity Crane suggests ‘is an expression of the aspectual shape of intentional states’ (p. 21)?

In this section, I will distinguish two stronger kinds of semantic equivalence relations. The first is that of having exactly the same intension, i.e. the same function from

⁴ It could fairly be pointed out that it is no less bold to take on the commitment that pains, or simple perceptual experiences, *can not* be directed upon impossible objects. So the present line of argument is not really an argument *for denying* Brentano’s thesis, on the present requirement for directedness, but an argument *against affirming* the thesis, on that requirement.

possible worlds to extensions. Notoriously, the terms ‘renate’ and ‘cordate’ are co-extensive but not co-intensive, whereas the predicates ‘is a renate’ and ‘is a renate and such that arithmetic is incomplete’ are both. Contexts susceptible to failure of substitution even for terms with the same intension are widely labelled ‘hyperintensional’. Reports of beliefs, desires and other propositional attitudes are hyperintensional, as exemplified by the non-equivalence of

(8) Jim believes Flukus is a cordate.

(9) Jim believes Flukus is a cordate and such that arithmetic is incomplete.

Modal statements are, by contrast, intensional but not hyperintensional. The same seems to go for reports of disposition. For example, the following are plausibly equivalent:

(10) The ball is disposed to attract an iron bar.

(11) The ball is disposed to attract a bar predominantly composed of stuff with atomic number 26.

This is plausibly explained by the co-intensiveness of ‘iron bar’ and ‘bar predominantly composed of stuff with atomic number 26’.

The second, even stronger, semantic equivalence relation I want to consider is that of having the same ‘Russellian meaning’. Roughly, two terms have same Russellian meaning iff they make the same contribution to the structured, Russellian proposition expressed by any sentence in which they figure.⁵ Somewhat more precisely, the two terms have the same Russellian meaning iff they have the same semantic structure, and the semantically atomic expressions that occupy the same place in the structure have the same semantic value, e.g. refer to the same object (if they are singular terms), or express the same property (if they are predicates), or denote the same truth-function (if they are

⁵ On the notion of a structured, Russellian proposition, see Salmon (1986) and King (2005).

connectives).⁶ For example, ‘ $\sqrt{4}$ ’ and ‘2’ have the same intension, but, as these terms differ in semantic structure, do not have the same Russellian meaning. The terms ‘Mark Twain’s father’ and ‘Samuel Clemens’s dad’ do by contrast have the same Russellian meaning, and therefore, plausibly, the same intension.

Contexts that are susceptible to substitution failure even for terms with the same Russellian meaning do not, as far as I know, have a special name in the semantics literature. I will therefore introduce the term ‘trans-Russellian’ for such contexts. Reports of beliefs, desires, or other propositional attitudes seem to be trans-Russellian. For example, it seems that (12) and (13) can differ in truth-value.

(12) Jim believes the man over there is Mark Twain’s father.

(13) Jim believes the man over there is Samuel Clemens’s dad.

Whether attitude reports are in fact trans-Russellian is, however, controversial, and something to which we shall return. On the other hand, it is hard to find any non-mental reports that are even plausible candidates for being trans-Russellian.

6. Strong and stronger substitution failures

We may now consider two necessary conditions for aspectual shape, corresponding to these two semantic equivalence relations and correlated kinds of substitution failure. One proposal is that states have aspectual shape only if reports thereof are susceptible to substitution failure for the terms that specify the objects of the state *even for* pairs of terms having the same intension. This requirement does not exclude beliefs and desires.

⁶ I simplify matters considerably here by disregarding context-sensitivity. A more precise definition would define sameness of Russellian meaning over ordered n-tuples of expressions, speakers, times, and other contextual factors relevant to reference determination.

In light of the apparent equivalence of such reports as (10) and (11) above, it might appear to exclude dispositions.

However, there are other, slightly more complex reports of a non-mental kind that are hyperintensional. Interestingly, and perhaps ironically, the reports in question are related to a form of words often used to help make explicit the aspectual shape of mental states. Suppose I am thinking of Jens Stoltenberg by having the thought that the Norwegian prime minister is about to resign while you are thinking of him by having the thought that the leader of the Norwegian Labour Party is about to resign. One way of trying to make explicit the different ways you and I are thinking of Jens Stoltenberg here – the different aspectual shapes of our thinking – is to say that you are thinking of him *as* the leader of the Norwegian Labour Party while I am thinking of him *as* the Norwegian Prime Minister.⁷ We noted above that such ‘as’ talk also finds natural application in certain simple physical situations: If *x* gravitationally attracts *y*, *x* attracts *y as* a thing with a mass; if *x* electrically attracts *y*, *x* attracts *y as* a thing having a certain charge. Very closely related to the latter are reports of the following form:

- (14) The ball attracts the bar with a force of approx $3.3 * 10^{-9}$ Newton in virtue of the fact that the ball has a mass of 50 kg, the bar a mass of 1 kg, and their distance apart is 1 m.

Such reports, however, create a hyperintensional context. Let ‘P’ abbreviate the sentence that follows ‘in virtue of the fact that’ in (14). Assuming the essentiality of

⁷ Woodruff Smith and McIntyre exemplify this use of ‘as’ constructions in the context of explaining Husserl’s view of the ‘noematic’ structure of acts of consciousness: ‘To do justice to Husserl, we ought to embellish the traditional gloss of intentionality: a consciousness is a consciousness of some particular thing under some particular description, i.e., *as* having certain properties (I see this object *as* a tree in bloom in Husserl’s garden)’ (1971: 550, emphasis added).

origin, 'P' has the same intension as 'Thorvald Stoltenberg fathered Jens Stoltenberg (if the latter exists) and P'.⁸ Interchanging these co-intensive expressions we get:

- (15) The ball attracts the bar with a force of approx $3.3 * 10^{-9}$ Newton in virtue of the fact Thorvald Stoltenberg fathered Jens Stoltenberg (if the latter exists) and P.

On a natural reading, (15) entails that the fact that Thorvald Stoltenberg fathered Jens Stoltenberg (if the latter exists) has something to do with the fact that the ball attracts the bar with the relevant force. But of course it does not. On a plausible reading, (14) lacks this false entailment. So, on a plausible reading, (14) and (15) are non-equivalent.

The states ascribed by such reports as (14) are not mental states. We have seen that such reports are susceptible to substitution failure for expressions that specify the objects upon which the states are directed even when the expressions are co-intensive. So these states meet the current requirement for having aspectual shape.

The obvious reaction to this problem is to shift to an even stronger necessary condition for aspectual shape: viz. the condition that states be ascribed by trans-Russellian reports. We saw above that beliefs, desires, and other propositional attitudes at least appear to meet this condition. No non-mental report I can think of even appears to meet it.⁹ So perhaps this is the way to sort the non-mental chaff from the mental wheat. The problem is that it is highly doubtful whether all mental states meet the requirement. Indeed, it is

⁸ If one doubts the essentiality of origin, one is free to substitute another necessary truth for the claim that TS fathered JS (if he exists).

⁹ This includes reports using the 'in virtue of' construction. From 'The soprano's tone made the glass break in virtue of its pitch and loudness' it seems to follow that 'The soprano's tone made the glass break in virtue of its frequency and amplitude', and vice versa. The example is adapted from Dretske (1989).

controversial whether such *paradigmatic* intentional states as beliefs and desires meet it since it is controversial whether attitude reports are trans-Russellian.¹⁰ When we turn to perceptual experiences and bodily sensations, it becomes yet more controversial that reports thereof are trans-Russellian. Tye (2000, 2003, 2006b) has recently argued, not implausibly, that even if standard attitude reports are trans-Russellian, standard reports of visual experiences and pains are not. To take the case of pain, let's assume it is true that

(16) Sam feels a pulsating pain in his left ankle.

Now, idealising ever so slightly, suppose that to throb is just the same as to pulsate, and that the ankle is the talocrural joint. On this supposition, many have the intuition that (16) is inter-deducible with

(17) Sam feels a throbbing pain in his left talocrural joint.

This intuition is certainly not without force.¹¹ So the current necessary condition for aspectual shape excludes non-mental states at the price of making it considerably more doubtful whether all mental states meet it. Again, it is telling that this necessary condition sets the standard of intentionality much higher than either Tye or Crane have even purported to attain in their arguments for the intentionality of bodily sensations and moods.

7. Conclusion

I conclude that no obvious variation over Crane's articulation of intentionality clearly escapes the trilemma that intentionality is either (i) too weak to exclude all non-mental

¹⁰ See Salmon (1986), Soames (1987), and Braun (1998) for representative criticism of the idea that attitude reports are trans-Russellian.

¹¹ See the Introduction, in particular sections 6.3 and 7.2, for further remarks on the corresponding intuition for reports of the form 'X looks F to S'.

states, (ii) too strong to include all mental states, or (iii) trivially sufficient for mentality. Crane articulation seems to me to be not only representative of standard interpretations of intentionality, but also to be among the most promising to Brentano's thesis. The obvious variations we have examined are also the natural ones. Of course, nothing of what we have said here conclusively refutes Brentano's thesis. It suggests, however, that the sufficiency of intentionality to mentality grounds a large and underappreciated challenge to the thesis.

Bibliography

Alter, T. & S. Walter (eds) (2007) *Phenomenal Concepts and Phenomenal Knowledge*.

Oxford: Oxford University Press.

Armstrong, D. (1968) *A Materialist Theory of the Mind*. London: Routledge.

Armstrong, D. (1978) *Universals and Scientific Realism*. Cambridge: Cambridge

University Press.

Austin, J. L. (1962) *Sense and Sensibilia*. Oxford: Oxford University Press.

Bigelow, J. & R. Pargetter (1990) 'Acquaintance with Qualia.' *Theoria*, 61: 129–47.

Block, N. (1995) 'On a Confusion about a Function of Consciousness' reprinted in N.

Block et al. (eds) *The Nature of Consciousness*. Cambridge, MA: MIT Press.

Block, N. (1998) 'Is Experiencing Just Representing?' *Philosophy and*

Phenomenological Research, 58: 663-670.

- Boghossian, P. & D. Velleman (1989) 'Colour as a Secondary Quality.' *Mind*, 98: 81–103.
- Boghossian, P. & D. Velleman (1991) 'Physicalist Theories of Color.' *Philosophical Review*, 100: 67–106.
- Botterill, G. & P. Carruthers (1999) *The Philosophy of Psychology*. Cambridge: Cambridge University Press.
- Braun, D. (1998) 'Understanding Belief Reports.' *Philosophical Review*, 105: 555–595.
- Braun, D. (2000). Russellianism and Psychological Generalisations. *Nous*, 34: 203–236.
- Braun, D. (2004) 'Review of Michael Thau, *Consciousness and Cognition*.' *Philosophy and Phenomenological Research*, 64: 484–491.
- Breckenridge, W. (2007) 'Against One Reason for Thinking that Visual Experiences have Representational Content.' *Philosophical Perspectives*, 21.
- Brentano, F. (1874) *Psychology from an Empirical Standpoint*. London: Routledge & Kegan Paul, 1973.
- Brewer, B. (1999) *Perception and Reason*. Oxford: Oxford University Press.

- Brewer, B. (2005) 'Perceptual Experience has Conceptual Content.' In E. Sosa & M. Steup (eds) *Contemporary Debates in Epistemology*. Oxford: Blackwell.
- Brewer, B. (2006) 'Perception and Content.' *European Journal of Philosophy*, 14: 165–181.
- Brewer, B. (2007) 'Perception and Its Objects.' *Philosophical Studies*, 132:87–97.
- Brewer, B. (forthcoming) 'How to Account for Illusion.' Ms available at URL=<...>. Resource last visited 22.05.07.
- Broad, C. D. (1923) *Scientific Thought*. London: Routledge & Kegan Paul.
- Burge, T. (1979) 'Individualism and the Mental' *Midwest Studies in Philosophy*, 4: 73–121.
- Byrne, A. (2001) 'Intentionalism Defended,' *Philosophical Review*, 110: 199–240.
- Byrne, A. (2005) 'Perception and Conceptual Content.' In E. Sosa & M. Steup (eds) *Contemporary Debates in Epistemology*. Oxford: Blackwell.
- Byrne, A. (2006) 'Color and the Mind-Body Problem.' *Dialectica*, 60: 223–44.
- Byrne, A. & D. Hilbert (1997) 'Colors and Reflectances.' In A. Byrne & D. Hilbert (eds) *Readings on Color. Vol 1*. Cambridge, MA: MIT Press.

- Campbell, J. (2002) *Reference and Consciousness*. Oxford: Oxford University Press.
- Campbell, K. (1981) 'The Metaphysic of Abstract Particulars.' *Midwest Studies in Philosophy* 6: 477–88. Page ref. to H. Mellor & A. Oliver (eds) *Properties*. Oxford: Oxford University Press, 1997.
- Chalmers, D. (2004) 'The Representational Character of Experience.' In B. Leiter (ed.) *The Future for Philosophy*. Oxford: Oxford University Press.
- Chalmers, D. (2006) 'Perception and the Fall from Eden.' In T. Gendler and J. Hawthorne (eds) *Perceptual Experience*. Oxford: Oxford University Press.
- Child, W. (1994) *Causality, Interpretation, and the Mind*. Oxford: Oxford University Press.
- Chisholm, R. (1957) *Perceiving: A Philosophical Study*. Ithaca, NY: Cornell University Press.
- Clark, A. (2000) *A Theory of Sentience*. Oxford: Oxford University Press.
- Cohen, A. (1999) *Think Generic! The Meaning and Use of Generic Sentences*. Stanford, CA: CSLI.

- Crane, T. (1992) 'The Nonconceptual Content of Experience.' In T. Crane (ed.) *The Contents of Experience*. Cambridge: Cambridge University Press.
- Crane, T. (1988) 'The Waterfall Illusion.' *Analysis*, 58: 142-7.
- Crane, T. (1998) 'Intentionality as the Mark of the Mental.' In A. O'Hear (ed.) *Current Issues in the Philosophy of Mind*. Cambridge: Cambridge University Press.
- Crane, T. (2001) *Elements of Mind*. Oxford: Oxford University Press.
- Cussins, A. (1990) 'The Connectionist Construction of Concepts.' In M. Boden (ed.) *The Philosophy of Artificial Intelligence*. Oxford: Oxford University Press.
- Dennett, D. (1991) *Consciousness Explained*. Boston: Little, Brown.
- Descartes, R (1971) *Meditations on First Philosophy*. In G. E. M. Anscombe & P. Geach (eds and trans.) *Descartes: Philosophical Writings*. Indianapolis: Boobs-Merrill.
- Dokic, J. & E. Pacherie (2001) 'Shades and Concepts.' *Analysis* 61: 193–202.
- Dretske, F. (1969) *Seeing and Knowing*. London: Routledge & Kegan Paul.
- Dretske, F. (1970) 'Epistemic Operators.' *Journal of Philosophy* 67: 1007–23

- Dretske, F. (1989) 'Reasons and Causes.' *Philosophical Perspectives*, 3: 1–15.
- Dummett, M. (1975) 'Wang's Paradox.' *Synthese*, 30: 301–324. Reprinted in his *Truth and Other Enigmas*. London: Duckworth, 1978.
- Evans, G. (1982) *The Varieties of Reference*. Oxford: Oxford University Press.
- Firth, R. (1949-50) 'Sense-Data and the Percept Theory.' *Mind*, 58-9. Reprinted in Schwartz, R. J. (ed.) *Perceiving, Sensing, and Knowing*. Garden City, NY: Doubleday, 1965.
- Fodor, J. (1983) *The Modularity of the Mind*. Cambridge, MA: MIT Press.
- Fodor, J. (1998) *Concepts. Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.
- Fodor, J. & Z. Pylyshyn (1988) 'Connectionism and Cognitive Architecture: A Critical Analysis.' *Cognition*, 28: 3-71. Page refs. to version available at URL=<<http://ruccs.rutgers.edu/ftp/pub/papers/jaf.pdf>>
- Forbes, G. (2006) *Attitude Problems*. Oxford: Oxford University Press.
- Foster, J. (2000) *The Nature of Perception*. Oxford: Oxford University Press.

- Frege, G. (1892) 'Über Sinn und Bedeutung.' *Zeitschrift für Philosophie und philosophische Kritik*, 100: 25–50. Trans. as 'On Sense and Meaning.' In P. Geach & M. Black (eds. & trans.) *Translations from the Philosophical Writings of Gottlob Frege*. Oxford: Blackwell, 1952.
- Goodman, N. (1951) *The Structure of Appearance*. (First edition) Cambridge, MA: Harvard University Press.
- Graff, D. (2001) 'Phenomenal Continua and the Sorites' *Mind*, 110: 905–935.
- Gregory, D. (2004) 'Imagining Possibilities', *Philosophy and Phenomenological Research*, 69: 327–348.
- Grice, P. (1961) 'The Causal Theory of Perception.' *Proceedings of the Aristotelian Society*, 35.
- Hanna, R. (2005) 'Kant and Nonconceptual Content.' *European Journal of Philosophy* 13: 247–290.
- Hardin, C. L. (1988) *Color for Philosophers. Unweaving the Rainbow*. Indianapolis: Hackett.
- Harman, G. (1990) 'The Intrinsic Quality of Experience' *Philosophical Perspectives*, 4. Page ref. to N. Block *et al* (eds) *The Nature of Consciousness*. Cambridge, Mass. MIT Press.

- Heck, R. (2000) 'Nonconceptual Content and the "Space of Reasons"' *Philosophical Review*, 109: 483–523.
- Higginbotham, J. (1983) 'The Logic of Perceptual Reports: An Extensional Alternative to Situation Semantics.' *Journal of Philosophy*, 80: 100–127.
- Jackson, F. (1977) *Perception: A Representative Theory*. Cambridge: Cambridge University Press.
- Jackson, F. (1982) 'Epiphenomenal Qualia' *Philosophical Quarterly* 32.
- Jackson, F. (1986) 'What Mary didn't Know.' *Journal of Philosophy*, 83: 291–295.
- Jackson, F. (2000) *Some Reflections on Representationalism*. Talk delivered at NYU, 25 April 2000. Ms available at URL=<<http://www.nyu.edu/gsas/dept/philo/courses/consciousness/papers/RepresentationalismNYU5April00.PDF>>
- Johnson, K (2004) 'On the Systematicity of Language and Thought.' *Journal of Philosophy*, 101: 111–39.
- Johnston, M. (2004) 'The Obscure Object of Hallucination.' *Philosophical Studies*, 103: 113–83.

- Kaiser, P. K. & R. M. Boynton (1996) *Human Color Vision*. Washington, DC: Optical Society of America.
- Kant, I. (1787) *Critique of Pure Reason*. Cited from N. Kemp-Smith (trans.) London: Pargrave-Macmillan, 2003.
- Kelley, D. (1980) 'The Specificity of Perception.' *Philosophy and Phenomenological Research*, 40: 401–5.
- Kelly, S. (2001) 'Demonstrative Concepts and Experience.' *Philosophical Review*, 110: 397–420.
- King, J. C. (2005) 'Structured Propositions.' *The Stanford Encyclopedia of Philosophy* (Summer 2005 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/sum2005/entries/propositions-structured/>>.
- Kuhn, T. S. (1962) *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Lerman, H. (ms.) 'Demonstrative Concepts and the Experience of Properties.'
- Lewis, D. (1980) 'Veridical Hallucination and Prosthetic Vision.' *Australasian Journal of Philosophy*, 58: 239–249.

- Lewis, D. (1983) 'Individuation by Acquaintance and by Stipulation.' *Philosophical Review*, 92: 3–32.
- Lewis, D. (1986) 'What Experience Teaches.' *Proceedings of the Russellian Society*, 13: 29–57.
- Loar, B. (1990/1977) 'Phenomenal States,' *Philosophical Perspectives* 4: 81–108. Reprinted with revisions in N. Block et al. (eds) *The Nature of Consciousness*. Cambridge, MA: MIT Press.
- Lycan, W. (1995) 'A Limited Defence of Phenomenal Information,' in Thomas Metzinger (ed.) *Conscious Experience*. Exeter: Imprint Academic.
- Lycan, W. (2003) 'Perspectival Representation and the Knowledge Argument,' in Q. Smith and A. Jokic (eds) *Consciousness. New Philosophical Perspectives*. Oxford: Oxford University Press.
- McDowell, J. (1979) 'Truth-Value Gaps.' Reprinted in McDowell (1998c).
- McDowell, J. (1982) 'Criteria, Defeasibility, and Knowledge' *Proceedings of the British Academy*, 68: 455–466. Page refs to reprinting in his *Meaning, Knowledge, and Reality*. Cambridge, MA: Harvard University Press, 1998.
- McDowell, J. (1984) '*De Re* Senses.' Reprinted in McDowell (1998c).

- McDowell, J. (1994) *Mind and World*. Cambridge, MA: Harvard University Press.
- McDowell, J. (1998a) 'Having the World in View: Sellars, Kant, and Intentionality.' *Journal of Philosophy*, 98: 431–490.
- McDowell, J. (1998b) 'Reply to Commentators.' *Philosophy and Phenomenological Research*, 58: 403–31.
- McDowell, J. (1998c) *Meaning, Knowledge, and Reality*. Cambridge, MA: MIT Press.
- McDowell, J. (forthcoming) 'Conceptual Capacities in Perception.' Preprint available at
URL = <[http://www.kcl.ac.uk/kis/schools/hums/philosophy/handouts/berlin
paper.pdf](http://www.kcl.ac.uk/kis/schools/hums/philosophy/handouts/berlin%20paper.pdf)>
- Mack, A & I. Rock (1998) *Inattentional Blindness*. Cambridge, MA: MIT Press.
- McLaughlin, B. (2003) 'Colour, Consciousness, and Colour Consciousness.' In Q. Smith & A. Jokic (eds) *Consciousness: New Philosophical Perspectives*. Oxford: Oxford University Press.
- Magidor, O (2007) *Category Mistakes*. Oxford University D. Phil. Thesis.
- Margolis, E. & S. Laurence (eds.) (1999) *Concepts: Core Readings*. Cambridge, MA: MIT Press.

- Martin, M. G. F. (1997) 'Sense, Reference and Selective Attention: The Shallows of the Mind.' *Proceedings of the Aristotelian Society. Supplementary Volume*. 71.
- Martin, M. G. F. (2002) 'The Transparency of Experience.' *Mind and Language*, 17: 376–425.
- Martin, M. G. F. (2004) 'The Limits of Self-Awareness.' *Philosophical Studies*, 103: 37–89.
- Moore, G. E. (1903) 'The Refutation of Idealism.' *Mind*, 12: 433–53. Page ref. to reprinting in his *Philosophical Studies*. London: Routledge and Kegan Paul, 1922.
- Moore, G. E. (1918–9) 'Some Judgements of Perception.' *Proceedings of the Aristotelian Society*, 19.
- Nemirow, L (1980) 'Review of *Mortal Questions* by Thomas Nagel.' *Philosophical Review*, 89: 473–477.
- Nes, A. (2006) *Content in Thought and Perception*. Oxford University D.Phil. thesis. Available at URL=<webpace.utexas.edu/deverj/personal/test/thoughtperception.pdf>
- Noë, A. (2004) *Action in Perception*. Cambridge, MA: MIT Press.

- Nozick, R. (1981) *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Palmer, S. (1999) *Vision Science: Photons to Phenomenology*. Cambridge, MA: MIT Press.
- Papineau, D. (2002) *Thinking About Consciousness*. Oxford: Oxford University Press.
- Papineau, D. (2007) 'Phenomenal and Perceptual Concepts.' In Alter & Walter (2007).
- Peacocke, C. (1983) *Sense and Content*. Oxford: Oxford University Press.
- Peacocke, C. (1984) 'Colour Concepts and Colour Experience.' *Synthese*, 58: 365–82.
- Peacocke, C. (1986) 'Analogue Content.' *Proceedings of the Aristotelian Society, Supplementary Volume*, 60: 1–17.
- Peacocke, C. (1989) 'Perceptual Content.' In J. Almog, J. Perry & H. Wettstein (eds) *Themes from Kaplan*. Oxford: Oxford University Press.
- Peacocke, C. (1992) *A Study of Concepts*. Cambridge, MA: MIT Press.
- Peacocke, C. (1994) 'Nonconceptual Content: Kinds, Rationals, and Relations.' *Mind and Language*, 9. 419–429. Page refs. to repr. in Gunther (2003).

- Peacocke, C. (1998) 'Nonconceptual Content Defended.' *Philosophy and Phenomenological Research*, 58: 381–388.
- Peacocke, C. (2001) 'Does Perception Have a Nonconceptual Content?' *Journal of Philosophy*, 98: 239–264.
- Pietrowski, P. & G. Rey (1995) 'When Other Things Aren't Equal. Saving Ceteris Paribus Laws from Vacuity.' *British Journal for the Philosophy of Science*,
- Pitcher, G. (1970) *A Theory of Perception*. Princeton, NJ: Princeton University Press.
- Pitt, D. (2004) 'The Phenomenology of Cognition, Or, What Is It Like to Think That P?' *Philosophy and Phenomenological Research*, 69: 1–36.
- Priest, G. (1987) *In Contradiction: A Study of the Transconsistent*. Dordrecht: Martinus Nijhoff.
- Pylyshyn, Z. (2003) *Seeing and Visualising. It's Not What You Think*. Cambridge, MA: MIT Press.
- Raffman, D. (1995) 'The Persistence of Phenomenology.' In T. Metzinger (ed.) *Conscious Experience*. Schoningh-Verlag Publishers.
- Raffman, D. (2000) 'Is Perceptual Indiscriminability Nontransitive?' *Philosophical Topics*, 28: 153–175.

- Rensink, R. A., J. K. O'Regan & J. J. Clark (1997) 'To See or Not to See: The Need for Attention to Perceive Changes in Scenes.' *Psychological Science* 8: 368–373.
- Richard, M. (1990) *Propositional Attitudes. An Essay on Thoughts and How we Ascribe Them*. Cambridge: Cambridge University Press.
- Robinson, W. (2005) 'Thoughts Without Distinctive Non-Imagistic Phenomenology.' *Philosophy and Phenomenological Research*, 70: 534–562.
- Russell, B. (1912) *The Problems of Philosophy*. Oxford: Oxford University Press.
- Salmon, N. (1986) *Frege's Puzzle*. Cambridge, Mass: MIT Press.
- Searle, J. (1983) *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.
- Sedivy, S. (1996) 'Must Conceptually Informed Perceptual Experience Involve Non-Conceptual Content?' *Canadian Journal of Philosophy*, 26: 413–431.
- Shoemaker, S. (2003) 'Content, Character, and Color.' *Philosophical Issues* 13: 253–278.

- Siegel, S. (2005) 'The Contents of Perception.' In E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy (Spring 2005 Edition)*, URL = <<http://plato.stanford.edu/archives/spr2005/entries/perception-contents/>>
- Siegel, S. (2006) 'What Properties are Represented in Perception?' In J. Hawthorne & T. Gendler (eds) *Perceptual Experience*. Oxford: Oxford University Press.
- Siewert, C. (2004) 'Is Experience Transparent?' *Philosophical Studies*, 117: 15–41.
- Snowdon, P. (1980–1) 'Perception, Vision, and Causation.' *Proceedings of the Aristotelian Society*, 81: 175–192.
- Soames, S. (1987) 'Direct Reference, Propositional Attitudes, and Semantic Content.' *Philosophical Topics* 15: 44–87.
- Soames, S. (1990). Belief and Mental Representation. In P. P. Hanson (ed.), *Information, Language, and Cognition*. Vancouver: University of British Columbia Press.
- Speaks, J. (2005) 'Is there a Problem about Nonconceptual Content?' *Philosophical Review*, 114: 359–398.
- Sperling, G. (1960) 'The Information Available in Brief Visual Presentations.' *Psychological Monographs*, 74: 1–29.

Stalnaker, R. (1984) *Inquiry*. Cambridge, MA: MIT Press.

Stalnaker, R. (1998) 'What Might Nonconceptual Content Be?' *Philosophical Issues*, 9: 339–352.

Stalnaker, R. (1999) *Context and Content*. Oxford: Oxford University Press.

Stoljar, D. (2004) 'The Argument from Diaphanousness', in M. Ezcurdia, R. Stainton and C. Viger (eds) *New Essays in the Philosophy of Language and Mind. Supplemental volume of The Canadian Journal of Philosophy*. Calgary: University of Calgary Press, pp.341–390. Page refs. to version available online at URL=<<http://philrsss.anu.edu.au/~dstoljar/onlinepapers/Diaphanousness.pdf>>

Thau, M. (2002) *Consciousness and Cognition*. Oxford: Oxford University Press.

Travis, C. (2004) 'The Silence of the Senses.' *Mind*, 113: 57–94.

Tye, M. (1984) 'The Adverbial Approach to Visual Experience.' *Philosophical Review*, 93: 195–225.

Tye, M. (1992) 'Visual Qualia and Visual Content.' In T. Crane (ed.) *The Content of Experience*. Cambridge: Cambridge University Press.

Tye, M. (1995) *Ten Problems of Consciousness*. Cambridge, Ma: MIT Press.

Tye, M. (2000) *Consciousness, Color and Content*. Cambridge, Ma: MIT Press.

Tye, M. (2002) 'To PANIC or Not to PANIC? – Reply to Byrne.' *A Field Guide to Philosophy of Mind* (Winter 2002) Symposium on Michael Tye's *Consciousness, Color and Content* : 28–33. URL= <<http://host.uniroma3.it/progetti/kant/field/tyesymp.pdf>>

Tye, M. (2003) *Consciousness and Persons*. Cambridge, MA: MIT Press.

Tye, M. (2006a) 'Nonconceptual Content, Richness, and Fineness of Grain.' In T. Gendler & J. Hawthorne (eds) *Perceptual Experience*. Oxford: Oxford University Press.

Tye, M. (2006b) 'The Thesis of Nonconceptual Content.' *European Review of Philosophy*, 6: 7–31.

Warnock, G. J. (1954–55) 'Seeing.' *Proceedings of the Aristotelian Society*, 60.

Weiskrantz, L. (1990) *Blindsight: A Case Study and Implications*. 2nd Edition. New York: Oxford University Press.

Williams, D. C. (1953) 'The Elements of Being' *Review of Metaphysics*, 7: 3–18. Page ref. to D. H. Mellor & A. Oliver (eds) *Properties*. Oxford: Oxford University Press, 1997.

- Williamson, T. (1990) *Identity and Discrimination*. Oxford: Blackwell.
- Williamson, T. (2000) *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Williamson, T. (2006) 'Conceptual Truth.' *Proceedings of the Aristotelian Society*.
Supplementary Volume, 80: 1–41.
- Williamson, T. (2007) 'On Being Justified in One's Head.' In M. Timmons, J. Greco &
A. R. Mele (eds) *Rationality and the Good: Critical Essays on the Ethics and
Epistemology of Robert Audi*. Oxford: Oxford University Press.
- Wittgenstein, L (1942) *The Blue Book*. Published in *The Blue and The Brown Books*.
Harper Perennial.
- Wittgenstein, L. (1977) *Remarks on Colour*. G.E.M. Anscombe (ed.), L. L. McAlister
and M. Schattle (trans.) Berkley: University of California Press.
- Woodruff Smith, D and R. McIntyre (1971) 'Intentionality Via Intensions.' *Journal of
Philosophy* 68 (18): 541–561.
- Wright, C. (1975) 'On the Coherence of Vague Predicates.' *Synthese* 30: 325–365.