

Getting Properties In Mind

by

Dominic Adrian Alford-Duguid

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy

Department of Philosophy
University of Toronto

© Copyright by Dominic Adrian Alford-Duguid 2016

Getting Properties In Mind

Dominic Adrian Alford-Duguid
Doctor of Philosophy
Department of Philosophy
University of Toronto
2016

Abstract

Perceiving a property—a colour, shape, or size—can permit us to think about it. For example, upon seeing a red round ball, you could think *that is red* or *it's that shape*. Yet philosophers devote less attention to perception-based thought about properties than to perception-based thought about objects like tables and chairs. As a result, few recognise how a right account of perception-based thought about properties has surprising (and sometimes shocking) consequences for perception and thought. This dissertation uncovers some of those consequences. I establish claims about perception-based thought about properties—about when it occurs, how it occurs, and why it matters. I then use these claims to argue for a series of controversial theses about perception and thought.

Acknowledgments

Imogen Dickie was my supervisor. She taught me how to be a philosopher, and why it matters to be a good one. I will forever be in her debt.

In addition to Imogen, I enjoyed the luxury of a large and fierce committee. A rigorous, creative, and nuanced interlocutor with an unparalleled range, Benj Hellie was an unfailingly insightful critic and guide. Gurpreet Rattan—a thinker of great subtlety—brought a patient and probing perspective to my work. Mohan Matthen taught me how to keep a weather eye on empirical research into perception, and his generosity was responsible for my initial exposure to cutting-edge philosophical work on perception. And in Diana Raffman I found a philosopher whose work balances simplicity, creativity, and precision: she makes it look easy; it isn’t.

I owe two further related debts. Phil Kremer and Bill Brewer served as internal and external examiner for my dissertation. I’m very grateful to Phil for his incisive and constructive feedback, especially my work lies outside his usual stomping grounds. Bill Brewer was an ideal external examiner: a major figure in the corner of philosophy we share, he read my work with great care, and his brilliant comments went right to the philosophical heart of my project. In addition, a careful reader will see his influence (this time through his writing) throughout the dissertation.

Now when I came to the Toronto graduate program I was told that my primary resource would be my fellow travellers in the program. But I couldn’t have imagined just how fabulous and brilliant a group these peers would prove to be. I’d like to thank, in no particular order: James Davies, Willie Costello, Manish Oza, Rima Basu, Emily Perry, Kevin Kuhl, Benjamin Wald, Robbie Howton, Adam Murray, Luke Roelofs, Jessica Wright, Lisa Doerksen, Dave Suarez, Damien Melamedoff, Diana Heney, Ken Boyd, Brian Embry, Rachel MacKinnon, Josh Brandt, Mark Fortney, Dan Rabinoff, Elliot Carter, Aaron Henry, Sean Michael Smith, Belinda Piercy, Mason Westfall, Johanna Thoma, Mike Arsenault, and (with extra special thanks) Zac Irving. Aside from

my fellow graduate students, I also learned a lot from conversation with Toronto faculty outside my committee. Among these, a few stand out: Rachel Barney, Brad Inwood, Andrew Sepielli, Jessica Wilson, and Nate Charlow. Thanks also to Ian Hacking for two singular conversations.

I was given the opportunity to visit the London philosophical scene for a semester, and it was one of the most stimulating and productive periods of my graduate career. Many of those I met left a lasting impression on my philosophical outlook. Among these, Ed Nettel, Henry Clarke, Maarten Steenhagen, Lea Salje, and Alex Geddes deserve special thanks. During this time I also had the chance to work with several London-area faculty. Bill Brewer, Ian Rumfitt, and Barry Smith were amazingly generous with their time and feedback, and all three left their mark on this project.

Outside of Toronto and London, I've also had occasion to meet many members of the wider philosophical community at conferences, workshops, and other venues. Of this group, I would especially like to thank (again, in no particular order): Wayne Wu, David Liebesman, Brian Rabern, Matthew Soteriou, James Stazicker, Kirk Ludwig, and John Campbell.

Beyond these colleagues from the philosophical world, this dissertation would not have been possible without the unbelievable support of my friends and extended family. For making me laugh, relax, and (most of all) see the many bigger pictures: Dina Fayerman, Trisha Jackson, Duncan Derry, Thomas Wong, Haig Bedrossian, Larissa Lay, Lily Tarba, Ed Rek, Dawn Woodward, Eleanor Burke, Alex Heeney, Evelyn Woodward-Rek, Subha Arulvarathan, Sarah Khan, A.J. Rowley, Paul Burke, Carol Off, Julia Rosenberg, Hilary Buttrick, Ian Radojkovic, Anne Mackenzie, Ann Bramson, Michael Poskanzer, Ethan Poskanzer, William Fortescue, Margaret Jacob, Clare Ungerson, Ilse Wong, Kathy Wazana, and Richard & Juliet Jung.

An eternal and heartfelt thanks also goes to my immediate family. My parents bravely follow their own independent dreams—as travellers, writers, adventurers, and human beings—yet have always lovingly and unhesitatingly sought to help their children do the same. My brother (and dear

friend) Tashi has been a source of support, wisdom, understanding, and care: an ally of the highest calibre.

My final and most loving thanks go to my partner and best friend Fatema Amijee. A fierce and incisive philosopher, as well as a caring and supportive companion, she has changed my life irrevocably since entering it. I would not be here without her; nor would I wish to be.

The dissertation is dedicated to the memory of my grandparents: Jack Alford, Anne Alford (*née* Hegewald), Robin Duguid (*née* Wearne), and Adrian Duguid. They are missed but never forgotten.

Table of Contents

3.2.3	Our capacity to ascribe special causal relations type-depends upon our capacity to ascribe properties as categorical.....	54
4	Conclusion.....	62
5	Looking Ahead: Perception and Thought.....	64
Chapter 3 Whence Come Properties Before The Mind?		67
1	Conceptual Topography.....	69
1.1	Perceiving the World	70
1.2	Varieties of Presentation	71
1.3	Phenomenal Content and Perceptual Presentation	74
2	How PREDICATIVE ATOMISM Constrains Perception.....	76
2.1	The Path from PREDICATIVE ATOMISM.....	76
2.2	Basic Cognitive Capacities and Immediate Justification	78
2.2.1	Our capacity to perceptually ascribe observable properties as categorical is cognitively basic.....	78
2.2.2	Classificatory Knowledge and Phenomenal Content.....	80
3	An Internalist Wedge.....	84
4	No Prospects for PHENOMENAL INTERNALISM	89
4.1	Edenic Sceptical Folly	89
4.2	Inelastic Metaphysics	91
5	Prospects for PHENOMENAL EXTERNALISM	94
6	Looking Back to Look Forward	95
Chapter 4 The Limits of Reference		97
1	The Challenge in Frege's Framework	98
1.1	Reference and Understanding	99
1.1.1	Frege on Sense and Reference	99
1.1.2	Two Aspects of Reference	101
1.2	Incomplete Senses.....	106

1.2.1	Composition of Fregean Thoughts.....	107
1.2.2	The Complete/Incomplete Distinction Within Language	109
1.2.3	How Frege Extends Incompleteness to Reference	110
1.2.4	How Frege Extends Incompleteness to Sense.....	111
1.3	No Need for Predicative Reference	114
1.3.1	An Appeal to the History of Metaphysics.....	114
1.3.2	An Argument from Fregean Austerity	116
1.3.3	An Argument from Classification	118
2	The Generality of the Challenge	119
2.1	The Challenge Generalised I—Mental Representation.....	120
2.1.1	On the Task of a Formal Semantic Theory	120
2.1.2	An Informative Objection.....	124
2.1.3	The Challenge for Mental Representation	127
2.2	The Challenge Generalised II—The Depth of the UNITY OF UNDERSTANDING PROBLEM	129
2.2.1	Russell Understood.....	130
2.2.2	Unity Without Structure	132
3	A Landscape of Responses	134
4	Conclusion.....	137
	Conclusion: Another Variety of Reference	139
	Bibliography	145

Introduction

Viewed from a distance, my dissertation is about a pervasive feature of the mind: its capacity to represent that things—objects, events, and other sundry particulars—are some way. These ‘ways’ are what metaphysicians have variously called ‘properties’, ‘features’, ‘modifications’, or ‘attributes’. We know them, at least in the first instance, as the colours, shapes, and sizes of what we see; the texture, density, and temperature of what we touch; and the sounds, smells, and flavours we hear, smell, and taste. Once we move beyond these primitive observable features of the world, the class of properties becomes both large and heterogeneous. For example, at the level of fundamental physics we speak of charm, spin, and other properties equally strange.

Viewed up close, I’m concerned with a fairly specific question: to what extent is the way of thinking about properties made available by perception analogous to the way of thinking about objects made available by perception? Where my dissertation departs from the well-trodden path defined by previous work on mental representation—both work within so-called ‘naturalised semantics’ (which explains mental representation wholly in terms of non-representational notions such as causation) and work within generally non-naturalist Fregean programmes (which explains mental representation in terms of representation-laden notions such as justification and rationality)—is to take this apparently innocuous question and push it hard.¹

The dissertation has four chapters and a conclusion. Chapter 1 uses an overlooked puzzle to re-fashion the debate over perceptual experience. It is an old observation that when thinking about objects on the basis of perception, a subject can make mistakes about an object’s features without losing her ability to think about the object. What few realise is that perception-based thought about properties is similarly tolerant of error. I argue that this tolerance should lead us to reject a view of

¹ Influential attempts to naturalise mental representation include Field (1978); Dretske (1981); Fodor (1987); Millikan (1989). Fregeans include Strawson (1959; 1974); Evans (1982); Peacocke (1992; 2008); Campbell (2006b).

perceptual illusion that has become pervasive. This view holds that just as we lie to others by saying something false, illusory experience misleads by misrepresenting how things stand in the world. Many accounts of perception carry commitment to this view of illusion. As a result, my argument also extends to undermine these accounts, and thus motivates the search for a radical alternative account of perception.

Chapter 2 concerns the structure of our thought about properties. During his Logical Atomist phase, Russell argued that our capacity for descriptive identification of objects (e.g. that the nearest ball is round) depends upon our capacity for perceptual demonstrative thought about objects. Chapter 2 develops a powerful new argument for a parallel view for thought about properties: the view I call ‘PREDICATIVE ATOMISM’. This is a version of the old empiricist claim that to think about unobservable properties (such as spin or charge), a subject must be able to think perception-based thoughts about observable properties (such as red or circular). My argument secures a deeper anchor for PREDICATIVE ATOMISM by exploiting neglected connections between thought about properties and our grasp of causation.

In Chapter 3 I use my argument for PREDICATIVE ATOMISM to undermine a range of accounts of perceptual experience. The accounts all presuppose a quite widespread view, which is known as ‘PHENOMENAL INTERNALISM’. It is a view about what determines the qualitative character of perceptual experience. Put intuitively, PHENOMENAL INTERNALISM says that what it is like to undergo a perceptual experience wholly depends upon facts about ‘what’s in the head’. Less intuitively, the view holds that subjects may have different objects and properties in their respective environments, and may differ in the relations they have to these entities, yet what it is like for these subjects when perceiving can remain the same.

My final chapter resurrects a neglected challenge from Dummett (1973/1981). His target was a claim about language: the Fregean thesis that just as names stand for their bearers—a relationship standardly called ‘reference’—predicates also stand for properties (e.g. ‘red’ stands for the property

of redness). Dummett's was an internal verificationist challenge to Frege's framework. In contrast, my version ditches Dummett's verificationism, and has a generality that enables it to undermine even highly non-Fregean attributions of reference to both predicative expressions and predicative concepts.

The dissertation concludes with my solution to the generalised Dummettian challenge. I use resources from Chapters 1 and 2 to show that perception-based thought about properties—thought of the kind that underwrites PREDICATIVE ATOMISM— involves a variety of reference to properties.

Chapter 1

Thinking Through Illusion

This chapter uses an overlooked puzzle to re-fashion the debate over perceptual experience. The puzzle concerns perception-based thought about properties. It parallels a more familiar puzzle that arises for perception-based thought about objects. This familiar puzzle takes off from the observation that perception can enable a subject to think about an object while at the same time misleading the subject as to what the object is like. My puzzle takes off from the parallel observation for the case of properties: perception of a property can enable thought about it, while at the same time misleading the subject as to what it is like.

The chapter has five parts. §1 introduces the puzzle. §2 sketches what has become a standard explanation of the relationship between visual experience and rational belief formation. §3 uses the puzzle to generate a new trilemma for the standard explanation. §4 argues that the right response is to jettison the standard explanation. §5 traces how this response re-shapes the debate over perceptual experience.

1 The Tolerance Puzzle for Perceptual Ascription

This section develops the puzzle I want to introduce. To get the puzzle in view, I begin with the more familiar parallel puzzle that arises for perception-based thought about objects.

Suppose a ball of red yarn rolls into view, followed by your curious cat. Because you've seen the ball, you're now in a position to form certain kinds of thoughts about it. We standardly express these thoughts with sentences like 'that is red', 'that is round', or 'that is rolling'. Philosophers call these thoughts—those made available only by a perceptual link with the object they are about—'perceptual demonstrative thoughts'. Now philosophers have long supposed that a subject can entertain a perceptual demonstrative thought about an object even when the beliefs she rationally

forms about the object on the basis of her perceptual link with it are almost all false.¹ For instance, I might see a large red chair, but mistake it for something small (perhaps because my eyesight has begun to fail). Yet these errors do not undermine my ability to think about the chair. I might still rationally form false beliefs about it. These beliefs I would express with sentences such as ‘that is red’, and ‘that is large’. Let’s call this mark of perceptual demonstrative thought ‘tolerance’.² The tolerance of perceptual demonstrative thought generates a familiar puzzle: how can perception mislead a subject about an object’s properties, yet still enable her to think about it? Call this the ‘tolerance puzzle for perceptual demonstrative thought’.

How tolerant is perceptual demonstrative thought? I’m not sure. Some theorists restrict tolerance by endorsing a version of sortalism about perceptual demonstrative thought (hereafter ‘sortalism’). ‘Sortalism’ is the view that a subject can think a perceptual demonstrative thought about an object only if she successfully classifies the object according to its kind.³ For instance, a sortalist might insist that to think *that is tall* of a tree requires recognising the object as a tree. Sortalism entails restrictions on the tolerance of perceptual demonstrative thought: a subject can think such a thought about a perceived object only if she can correctly isolate properties of the object whose identification enables classification of the object according to its kind. The more specific the kind a subject must use when classifying an object for the purposes of perceptual demonstrative thought, the fewer of an object’s properties she can get wrong without undermining her capacity to think

¹ Cf. Evans (1982); Campbell (2002); Recanati (2012); Dickie (2015).

² My use of ‘tolerance’ is not borrowed from the literature on linguistic vagueness. Since Wright (1975, pp. 333–4), philosophers of language have used ‘tolerance’ to describe a particular feature of vague predicates. As Wright puts it, when ‘there are degrees of change in point of size, maturity and colour which are insufficient to alter the justice with which some specific predicate of size, maturity or colour is applied’, those predicates are ‘tolerant’. (More recently, Raffman (2014, §5.6) spells out a notion of tolerance that applies to vague predicates without generating paradox.) In contrast, perceptual demonstrative thought can be tolerant *in my sense* even if the metasemantic predicate ‘is about object o’ is not tolerant *in Wright’s sense* when applied to thoughts formed in response to increasingly misleading experience.

³ A sortalist must explain, *inter alia*, which kinds (of the many an object may fall under) a subject must use when classifying an object for the purpose of thinking about it. Sortalists include Strawson (1959), Quine (1950), Dummett (1973/1981), and Wiggins (1997). Dickie (2014) defends a less extreme form of sortalism. For arguments against sortalism, see Ayers (1974; 1997), Campbell (2002; 2006a), and Goodman (2012).

about it. However even a demanding version of sortalism will not make perceptual demonstrative thought wholly intolerant.

Now let's shift our focus from perception of and thought about objects, to perception of and thought about observable properties. When the red yarn ball rolls into view, I'm also able to think about its colour, its shape, its size, and its state of motion. I might express these thoughts with sentences like 'the ball is that colour', 'it's that shape', and (in the pure case) 'it's like this'. We'll say that to think about the object's properties in this way—a way made available by my perceptual link with the object's properties—is to 'perceptually ascribe' them. Perceptual ascription of properties thus functions as the predicative analogue of perceptual demonstrative identification of objects.

What nobody seems to have recognised is that perceptual ascription generates its own tolerance puzzle: how can perception mislead a subject about a property, yet still enable her to perceptually ascribe it? Attention to this overlooked puzzle transforms the debate over perceptual experience.

To bring out the tolerance puzzle for perceptual ascription, we must show that a subject (a) can be rationally misled about a property on the basis of her visual experience, yet (b) nevertheless perceptually ascribe the property. I'll take these each in turn.

We use properties to group or classify objects. Thus when presented with a collection of objects, and asked to group them by shape, we group like with like: the square objects in one group; the round objects in another; and so on. When perception misleads us about an object, and we're asked to group the object with others like it, we misclassify it: if a round object looks oval, we group it with oval objects; if a red object looks orange, we group it with orange objects; and so on.

A parallel story applies to properties. We use *higher-order* properties to classify the properties of objects. For example, if we sort a collection of coloured lights according to their hues (hue being

a higher-order property of colours)⁴, we group lights together just in case their colours have the same hue. As a result, when perception misleads us about hue, we misclassify colours. If a red object looks orange, we classify it with the other oranges, not with the reds.

Consider a case where you perceive three objects, each of which instantiates a different shade of red. (For those chary of drawing conclusions on the basis of colour cases, do not fear! I talk about shape in a moment.) Suppose you can tell by looking that these shades are distinct. So far nothing precludes your being able to perceptually ascribe each of the shades. If it turns out that you make a (rationally blameless) mistake when asked to order the shades from brightest to darkest, your visual experience will count as misleading you about which shades these objects possess. Would this classificatory mistake undermine your capacity to perceptually ascribe the shades? I think not. You're able to tell the shades apart, and that seems sufficient to permit successful perceptual ascription. To ask more of you—in this case an infallible ability to order shades by comparative brightness—would render perceptual ascription too difficult. Indeed it could render perceptual ascription of an object's colour nearly impossible.

The three-colour case involves misleading perception of three colour properties. Common sense suggests that this misleading perception nevertheless permits perceptual ascription of the colours. We could (and indeed I shall in a moment) construct a parallel case for shape properties. So we have candidate cases that clearly satisfy (a)—a subject is rationally misled about a property on the basis of her visual experience—and that seem intuitively to satisfy (b)—the subject can perceptually ascribe the property about which she is misled. In order to firmly establish the tolerance of perceptual ascription, though, we need more than mere intuitive support for (b).

⁴ Compare Hardin (1988, p. 25): ‘We may distinguish three dimensions of perceived color: hue, brightness, and saturation... the hue of a color is its [degree of] redness, or greenness, or yellowness, or blueness. White and black and the grays are the colors with zero hue...’

Additional support comes from a further parallel between perceptual demonstrative thought about objects and the perceptual ascription of properties. Consider an ordinary case of perceptual demonstrative thought. You come across an object. Your perceptual attention focusses on it. This attentional focus results in an influx of perceptual information about the object—where it is, what properties it possesses, what changes it undergoes.⁵ Much of this information you take at face value when forming beliefs. This growing collection of beliefs evolves in response to changes in the nature and structure of the incoming information.

When do these perceptual demonstrative beliefs—which we'd express using sentences like ‘that is red’ or ‘that is round’—count as being about the object to which you attend? An old and attractive answer is that your perceptual link with the object must govern your dispositions to form and maintain these perceptual demonstrative beliefs.⁶ These dispositions must organise your beliefs in response to a perceptual link with the object they purport to be about. They must also preserve a degree of internal rational coherence (e.g. no obviously inconsistent beliefs permitted).⁷ Most importantly, these dispositions must treat the perceptual link as *authoritative*: incoming perceptual information will usually trump other putative sources of information about an object.

Note that the previous paragraph addresses only the question of *when* perceptual demonstrative thoughts count as being about an object. It leaves open the much harder question of

⁵ For a textbook account of visual attention to objects, including a discussion of how attention responds to changes in the attended object, see Palmer (1999, pp. 554–6; 561–3).

⁶ These dispositions constitute our perception-based governing or ‘controlling conception’ of an object (cf. Evans 1982, Ch. 5). The recent literature on mental files and singular thought descends from Evans's discussion, and includes specific proposals about the constraints that a governing conception must satisfy for a body of beliefs to be about a particular object (cf. Dickie 2010b; Recanati 2012).

⁷ Compare Evans (1982, p. 146): ‘demonstrative thoughts take place in the context of a *continuing* informational link between subject and object: the subject has an evolving conception of the object, and is so situated *vis-à-vis* the object that the conception which controls his thinking is disposed to evolve according to changes in the information he receives from the object.’

how a perceptual link with an object does its aboutness-fixing work. Fortunately, nothing I want to say in this chapter requires taking a stand on how this harder question should be answered.⁸

I now turn from perception-based thought about objects to perception-based thought about properties. Whatever theorists eventually say about *how* aboutness-fixing works for these two types of thought, the stories about *when* it works remain almost exactly the same.

Consider a subject whose visual experience misleads him about the shape of a round vase, so that he mistakenly takes the vase to be oval. (I follow common usage, here and throughout, in assuming that being oval and being round are incompatible properties.) His perception-based beliefs about the vase's shape will be formed in response to information delivered by his attentional link with the shape.⁹ The dispositions that cause him to form beliefs in this fashion not only treat his perceptual link to the shape as authoritative, but also seek to preserve a degree of internal coherence amongst the resulting beliefs about the vase's shape (these beliefs will be about the shape's higher order properties—a departure from the perceptual demonstrative case: those beliefs were about properties of the perceived object). Suppose our subject sees that the vase rolls, and as a result comes to believe that *that shape* endows objects with the ability to roll (assume he already knows that an object's capacity to roll depends in part its shape). Now suppose that the vase stops rolling, and the situation evolves so that the subject would no longer classify the vase among the oval things (so the perceptual experience ceases to be misleading). In that case the subject will retain his belief that *that shape* endows objects with the ability to roll—though now he will revise his belief about which shape *that shape* happens to be (namely round, not oval).

⁸ The harder question remains a topic of ongoing debate. Cf. Evans (1982); Campbell (2002); Smithies (2011); Recanati (2012); Dickie (2015).

⁹ Carrasco (2011) provides an overview of the empirical literature on visual attention—including a survey of work on attention to observable properties (what psychologists call ‘feature-based attention’). While the focus of empirical investigation into feature-based attention diverges from that into object-based attention—for instance there is no analogue for feature-based attention of the literature on attentional object tracking (cf. Scholl 2001)—psychologists have come to recognise the existence and pervasiveness of both phenomena. For a philosophical overview of recent work on attention, see Wu (2014).

This pattern of belief revision reflects an attempt to preserve the same kind of internal coherence that we try to maintain when forming beliefs about objects on the basis of a perceptual link. Consider the object case first. ‘I thought it was large and far away, but now I see that it is small and nearby’ we say; ‘I thought the thing was red, given how it looked then, but now I see it is orange’. In cases of perceptual demonstrative thought about objects, our beliefs evolve to maintain rationality in response to information delivered through the attentional perceptual link that is underpinning perceptual demonstrative aboutness. The case I have just rehearsed demonstrates the same pattern for perceptual ascription. Suppose the subject is asked to make sense of his evolving beliefs about the object’s shape. ‘I saw that objects with that shape roll,’ he will say. ‘I simply mistook that shape for oval (given how the shape looked to me before), whereas now I see it is round.’ The subject’s beliefs have evolved to maintain rationality—the rationality of his beliefs about the shape—in response to the information that his perceptual link with the shape delivers.

Just as some endorse sortalism about perceptual demonstrative thought, one could endorse a kind of sortalism about perceptual ascription. On this type of sortalism, to perceptually ascribe an observable property a subject must classify the property according to one of its kinds: it is a shade, a shape, a height, etc.¹⁰ Hence if we’re sortalists about perceptual ascription, a thought we might express with ‘it’s like this’ must equally well be expressed by a sentence like ‘it’s that colour’ or ‘it’s that shape’ (compare how a traditional sortalist about perceptual demonstrative thought might insist that ‘that is red’ is always elliptical for a sentence such as ‘that table is red’ or ‘that house is red’).

I believe some type of sortalism about perceptual ascription *must* be true. Perception cannot lead a subject to mistake a round object for something triangular, yet at the same time permit him to

¹⁰ Cf. Peacocke (2008).

perceptually ascribe the object's actual (round) shape.¹¹ But just as with sortalism about perceptual demonstrative thought, sortalism about perceptual ascription is compatible with a degree of tolerance. To get a property's kind right—to correctly classify it as a shade, a shape, etc.—we need not know every observable higher-order property of the property.

So I take it that there are non-trivial limits on the tolerance of perceptual ascription: certain errors about the higher-order properties of observable properties—those errors that undermine our capacity to classify these properties according to their kinds—would prevent perceptual ascription of the properties. We can think about these limits in terms of determinates and determinables.¹² Given a determinate property—say scarlet—there may exist maximally dissimilar determinates such that if a perceptual experience were to lead someone to mistake the former for any of the latter, the experience would still permit the subject to perceptually ascribe the original property. These determinates—along with every determinate in between—would fall under a common minimally specific determinable property. Hence limits on tolerance would correspond to such determinable properties.

This way of thinking about limits on tolerance generates further questions. How determinate (or ‘narrow’) are these limit-defining determinables? How might these determinables differ *within* property categories (e.g. might perceptual ascription of certain determinate colours admit greater tolerance than ascription of others)? And how might they differ *between* property categories? For example, shape and colour plausibly afford different levels of tolerance: while it seems plausible that we can mistake a red object for purple while still thinking ‘that shade’ thoughts about the redness, it seems much less plausible that we can mistake a round object for something triangular while still

¹¹ Brewer (2011, p. 73) argues that there are limits to the kinds of ways an object can look—and thus limits to perceptual illusion. If he's right, it might be that a round object cannot look triangular. Yet even within these sorts of boundaries on illusion, it seems quite plausible that there are cases that go beyond the limits of tolerance.

¹² Armstrong (1997, §4.12) provides a canonical statement of the distinction between determinates and determinables. He takes his lead from Johnston (1921). For a recent influential discussion of the distinction, see Funkhouser (2006).

thinking ‘that shape’ thoughts about the thing’s actual round shape. It could also be that there isn’t a separate limit-defining determinable for each observable determinate. There might instead be a restricted class of privileged limit-defining determinables. Finally, apart from these abstract questions about the structure of tolerance, there remain hard questions about *why* these boundaries exist where they do.¹³

Despite the evident interest of these fine-grained questions about the tolerance of perceptual ascription, my primary focus in this chapter lies elsewhere. I shall argue that the tolerance of perceptual ascription leads us inexorably to reject a highly influential way of thinking about perceptual experience.

2 Common Sense View and the Tempting Explanation

My target will be a standard explanation of what I dub the ‘common sense view’ of the relationship between visual experience and rational belief formation. This section sketches the common sense view and introduces the standard explanation.

The common sense view captures three points of widespread agreement over the relationship between visual experience and rational belief formation:

1. Visual experience of an object can lead us to rationally form beliefs about the object, and these beliefs may be true.

For instance, a subject might see an oval vase and in response come to believe *that is oval*. And lo, the vase is oval, so her belief is true.

¹³ A prime place to look would be a right account of perceptual ascription. For example, Brewer (2005) argues (in response to an argument in Heck 2000) that our capacity to keep track of a property despite variations in perceptual circumstances is part of what enables us to think about it. Depending on how strongly we read this ‘tracking’ requirement, we will get different restrictions on the tolerance of perceptual ascription. Subjects are able to keep track of perceived properties through shifts in perspective, illumination conditions, etc. But how much of this tracking ability remains once perception begins to mislead us about observable properties?

2. Experience can also mislead us: the beliefs we rationally form about a perceived object may be false.

These kinds of cases are familiar. For instance, a subject might see her round vase through a distorting medium, and in response come to believe *that is oval*. But lo, the vase is round, not oval, so the belief is false.

3. Only some beliefs about an object formed in response to a visual experience of the object are rationally formed.

Strange are the cases that support 3. For instance, a subject might see a round vase, and the vase might *look* round to him, and yet somehow the subject might be caused by his experience to believe *that is oval*. But lo, the vase is round, not oval, so his belief is false. Though the belief is caused by the experience, absent further exculpatory information we cannot make sense of it as a rational response to the experience—the subject sees a vase that both is round and looks round, so how can his belief that it is oval, formed in response to his experience, count as rational?¹⁴

There is also a fourth kind of case that I have not mentioned in the presentation of the common sense view. As I develop it, the view remains silent about the rational significance of hallucinatory experience—bizarre experiences that lack an object (e.g. it seems to Macbeth as though a dagger floats before him, but lo, there is in fact no dagger there). My own view is that common sense delivers no clear verdict about hallucinatory cases.¹⁵

The combination of 1–3 constitutes the common sense view: only some beliefs formed in response to experience are rationally formed, and these may be either true or false.

¹⁴ Cf. Hellie (2011, p. 130; 2014, §3); Campbell and Cassam (2014, pp. 86–7).

¹⁵ In this respect I differ both from those like Matthen (2010) who insist we can rationally form beliefs in response to hallucinatory experience, and from those like Campbell (2002) who deny that these beliefs can be rationally formed.

The common sense view requires explanation. Philosophers find one explanation particularly tempting. Call this the ‘tempting explanation’.¹⁶ The tempting explanation appeals to a distinction between veridical and illusory visual experience. Put crudely, the distinction holds that veridical experience of an object gets the object’s properties right, while illusory experience of an object gets its properties wrong. Armed with this distinction, the tempting explanation explains the common sense view as follows:

- A. Whether a belief with a given content counts as a rational response to visual experience depends upon which content the experience possesses (where the ‘content’ of an experience is the condition under which it would count as accurate).¹⁷
- B. When the experience is veridical (i.e. when the content is accurate) an appropriately formed¹⁸ belief will be both true and rational.
- C. But when the experience is illusory (i.e. when the content is inaccurate) a belief with the same content as the first, and formed in the same way, will be false but rational.

Note that because it explains the common sense view in terms of the distinction between veridical and illusory experience, the tempting explanation treats the distinction between veridical and illusory experience as prior in order of explanation to the combination of claims about perception and rational belief formation that make up the common sense view.

I shall argue in §3 that the combination of the tempting explanation and the common sense view is *unstable*. This is a surprising—perhaps even shocking—result, because the tempting

¹⁶ Compare Byrne (2009, p. 429): ‘Perceptual experiences [have been conceived] as fallible, testifying (sometimes wrongly) about the subject’s external environment.’

¹⁷ Cf. Peacocke (1992, Ch. 3) for a seminal treatment of perceptual content in terms of accuracy conditions.

¹⁸ There remains room for disagreement over the question of what it takes for a belief to be ‘appropriately formed’ in response to experience (though nothing I say in this chapter depends upon how proponents of the tempting explanation settle such disagreements). For example, one might hold that a subject’s belief is appropriately formed in response to an experience only if it shares its content with the experience; others will add further conditions—such as the requirement that the belief be caused in a particular way by the experience.

explanation appears to be a very plausible explanation of the common sense view.¹⁹ One way to emphasise its plausibility is by way of a parallel with an attractive account of why beliefs formed in response to testimony can be false yet rational. When I testify, I use a sentence in a shared language to express something I wish to say. In a case of ordinary assertion, where I make my utterance with the aim of expressing a belief that I want my audience to share, successful testimony requires that my audience hear, understand, and believe what I say. If what I say is true, successful testimony will result in my audience forming a true belief. But if what I say is false, my audience will form a false belief. In both cases, what I say is the same, and my audience has no reason to doubt my word. These two facts explain why it is rational for my audience to form beliefs with the same contents in both cases.²⁰ What I say when I testify also constrains the sorts of beliefs my audience can rationally form on the basis of my testimony. For example, if I say ‘London is cold tonight’, my audience cannot, without substantial additional stage setting, rationally believe that London is warm tonight on the basis of my testimony.

In the testimony case, what I say acts as an intermediary between speaker and hearer. It explains my audience’s false (or true) beliefs in a way that preserves their rational status, and it also constrains the class of beliefs that it would be rational for my audience to form in response to my testimony. The seemingly tight parallel between what we’re explaining in the testimony case, on the one hand, and the common sense view, on the other, suggests that an explanation of the latter must identify a feature of visual experience that accomplishes the same explanatory task as the testimonial intermediary (what I say when I testify). Perceptual content is the best (if not the only) candidate.

¹⁹ For the view that we must explain rational restrictions in terms of *some* notion of perceptual content, compare Lewis (1994), McDowell (1994), Pryor (2000), and Hellie (2014); opposing views include Campbell (2002), Brewer (2011), and Dickie (2015). Those that explain visual phenomenology (including the veridical/illusory distinction) in terms of content that can sometimes be inaccurate—often called *intentionalists*—include Tye (2000), Byrne (2001; 2009), Chalmers (2004; 2006a), and Siegel (2010).

²⁰ Many deny that the rationality of testimonial uptake merely requires the absence of a positive reason to doubt a speaker. I’m not committed to the view here, but for persuasive defences see Burge (1993) and Dummett (1994).

Many will recognise their own views as falling under the broad umbrella of the tempting explanation. But in §§3–4 I shall use the tolerance puzzle from §1 to argue that we *cannot* appeal to perceptual content to explain the common sense view, if these contents are put to the same work as the testimonial intermediaries in the parallel case of misleading testimony. Abandoning the tempting explanation will force a radical re-think of some longstanding assumptions about perceptual experience. What we end up with in §5 is not merely a demand for a new account of rational belief formation, but an argument for jettisoning a deeply-rooted assumption about how to draw the distinction between veridical and illusory experience.

3 Tolerance Fork

In this section I raise a new and powerful trilemma for theorists of perceptual experience—what I shall call the ‘Tolerance Fork’.

Let’s stipulate that a subject stands in ‘perceptual contact’ with an object (or a property) iff she has a perceptual link with the object that sustains demonstrative thought about it (or a perceptual link with the property that sustains perceptual ascription of it).

The first step on the road to the trilemma is a principle connecting perceptual contact and perceptual content:

Contact and Content: Where an experience of o places the subject in perceptual contact with F , the experience’s content is accurate only if o is F .

Those already committed to the tempting explanation should accept *Contact and Content*. I shall argue in §4.2 that a proponent of the tempting explanation who rejects *Contact and Content* cannot accommodate the tolerance of perceptual ascription. In the meantime, a simple line of thought reveals why someone who accepts the tempting explanation is hard-pressed to reject *Contact and Content*. The tempting explanation chains perceptual content to how the world looks to a subject. When I see a cup on my table, for example, the cup will look red iff redness figures in the accuracy

conditions of my experience (i.e. the experience is accurate only if the cup is red). As a result, when combined with the tempting explanation, *Contact and Content* becomes the claim that a subject can perceptually ascribe only those properties that figure in how things look to her. And this claim appears well motivated. For instance, it helps explain why we cannot perceptually ascribe microscopic properties (e.g. mass or charge). It also explains why we cannot perceptually ascribe observable properties present in a pitch-black room.

The tolerance of perceptual ascription entails that a visual experience might mislead a subject with respect to an observable property, yet nevertheless permit her to perceptually ascribe the property about which she is misled. We've already established that someone—let's call him 'Gareth'—who sees a round vase but rationally believes it to be oval (because the circumstances of his perception make it look oval) satisfies these conditions. If asked, he would group the vase with oval objects, and so qualify as misled about the higher-order properties of the vase's shape (and thus about whether the vase is oval). At the same time, his dispositions to update and maintain his beliefs about the vase's shape are governed by a perceptual link with the shape—a fact that provides grounds for assuming that he could perceptually ascribe this shape.

Since Gareth's qualifies as a tolerance case, his visual experience puts him into perceptual contact with the vase's actual shape (roundness). *Contact and Content* uses this perceptual contact to import the shape of the vase into the content of Gareth's experience. Therefore the content of Gareth's experience is accurate only if the vase is round.

Yet in response to his misleading visual experience, Gareth also rationally believes that the vase is oval. The tempting explanation explains the rational status of this belief by insisting that the content of his experience is accurate only if the vase is oval.

Hence the content of Gareth's experience is accurate only if the vase is both round and oval—an apparently absurd result. Roundness and ovalness are incompatible shape properties. No object can possess both simultaneously.

Let's now provide a full dress version of the argument (note: a parallel argument could be run for any tolerance case):

1. *Contact and Content*: Where an experience of o places the subject in perceptual contact with F , the experience's content is accurate only if o is F .

Given **1** and the fact that Gareth's is a tolerance case,

2. The content of Gareth's experience is accurate only if his vase is round.

But Gareth's being a tolerance case also entails that

3. It would also be rational for Gareth to believe (in response to his experience) that his vase is oval.

The tempting explanation now makes an appearance:

4. *Tempting Explanation*: where it is rational to believe that o is F in response to an experience, the content of that experience is accurate only if o is F .

3 and **4** entail **5**:

5. The content of Gareth's experience is accurate only if his vase is oval.

Yet **2** and **5** are in tension: were both true, Gareth's experience would be accurate only if his vase is both oval and round. And that seems absurd.

Given **1–5**, a theorist of perceptual experience who accepts the tolerance of perceptual ascription faces a hard choice: somehow resolve the tension between **2** and **5**; reject *Contact and Content*; or reject the tempting explanation. This trilemma is the Tolerance Fork.

4 Two Alternative Replies to the Tolerance Fork

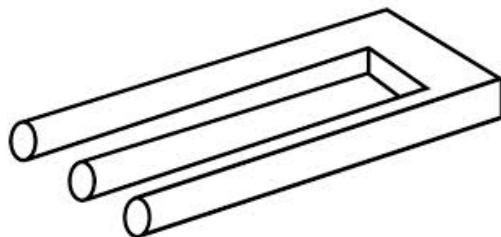
Before developing my own reply in §5—which is to reject the tempting explanation—I shall consider the other two potential replies to the Tolerance Fork. One attempts to defuse the tension between **2** and **5**. The other rejects *Contact and Content*. I shall argue that a proponent of the tempting explanation cannot accept either of them.

4.1 Defuse the Tension

There are two ways one might attempt to defuse the tension between **2** and **5**. The first embraces the possibility of experiences with inconsistent content. The second avoids positing inconsistent contents by insisting that **2** and **5** apply to contents of different types. I'll consider each of these options in turn.

4.1.1 Admit Inconsistent Perceptual Content

The possibility that visual experience might possess inconsistent perceptual content is familiar from debates over perception of ‘impossible objects’. When we see an impossible object our visual experience presents something that appears to possess incompatible properties. In the case of the devil’s pitchfork, for example, we’re inclined to say that a single object looks to have both two prongs and three:



Some philosophers insist that a right account of perceptual experience cannot permit inconsistent perceptual contents.²¹ If they are right, this first reply to the Tolerance Fork will not get off the ground.

Yet we need not side with these philosophers to undermine this first reply. For even if we allow that a visual experience can have inconsistent perceptual content, the proponent of the tempting explanation still lacks an answer to the Tolerance Fork. To see why, recall that according to the common sense view two experiences may lead subjects to rationally form beliefs with the same

²¹ Cf. Crane (1988).

contents, despite one belief being true and the other false. Now consider Sameen, a non-misled analogue of Gareth. Whereas Gareth sees a round vase and rationally believes it to be oval, Sameen sees an oval vase and rationally believes it to be oval. The two see differently shaped objects, yet their beliefs about the shapes have the same content, and both beliefs are rationally formed (though only Sameen's is true).

To capture these symmetries between Gareth and Sameen, the tempting explanation makes two moves. First: Sameen's belief is true, while Gareth's false, because Sameen's experience is accurate and Gareth's inaccurate (i.e. veridical and illusory, respectively). Second: Sameen and Gareth rationally form beliefs with the same content, despite the difference in their perceived surroundings, because their experiences share a common content.

Proponents of the tempting explanation lose this explanation of the symmetries between Sameen and Gareth if they insist that Gareth's experience possesses inconsistent content. For the explanation to work, Gareth and Sameen must have experiences with the same content. Thus if Gareth's experience has inconsistent content, so does Sameen's. Yet Sameen's experience must also be veridical, which requires that the content of her experience be capable of being accurate. But inconsistent contents cannot be accurate. As a result, to accept inconsistent content as an answer to the Tolerance Fork is to abandon the tempting explanation of the common sense view.

4.1.2 Multiply Types of Perceptual Content

A defender of the tempting explanation may insist that perceptual experiences admit different dimensions of accuracy—or (what comes to the same thing) different types of perceptual content. Admitting different types of perceptual content avoids the Tolerance Fork without incurring commitment to inconsistent content. Gareth's perceptual experience could (without inconsistency) include one type of content whose accuracy depends upon the perceived object being oval, and another type of content whose accuracy depends upon the perceived object being round.

As attractive as this line of response might seem, it fails. In order to drive this reply home an opponent must admit at least two types of content: α -content and β -content. α -content will explain why it is rational for Gareth to believe his vase is oval; β -content will explain why it is rational for him to believe that it's *that shape* (round).

In order to sustain the tempting explanation, this division between α - and β -content must be explanatorily prior to facts about which beliefs it would be rational to form in response to experience (facts that help compose the common sense view). My original reconstruction of the tempting explanation secured this explanatory priority by appeal to a notion of content designed to capture the way a perceived object looks from a subject's point of view. A round object may look round, or it may look oval (if the latter, we classify our experience as illusory). To capture these 'looks' facts—facts that concern perceptual phenomenology rather than rational belief formation—the tempting explanation introduces a notion of content: when a round object looks round, the content of my experience is accurate; when the same object looks oval, the content of my experience is inaccurate. The explanatory power of the tempting explanation comes from the fact that it derives this notion of content from facts about perceptual phenomenology, and then uses it to explain what it would be rational for a subject to believe in response to experience.

The tempting explanation thus faces a challenge. In order to distinguish α -content and β -content, the division must have a basis that is explanatorily prior to facts about which beliefs it would be rational to form in response to visual experience. Yet perceptual phenomenology offers no hope: objects do not possess the requisite conflicting looks—Gareth's vase doesn't look both round and oval. Since the vase *does* look oval, the tempting explanation can still use α -content to explain why it is rational for Gareth to believe it's oval. So the challenge for the tempting explanation is to secure β -content—a type of content that explains why it is rational for Gareth to believe his vase has

that shape—without explaining β -content in terms of facts about what it would be rational to believe in response to visual experience.

Proponents of the tempting explanation must look beyond perceptual phenomenology for resources to explain β -content. An attractive move is to explain β -content in terms of perceptual contact. A property would belong in the β -content of an experience iff the experience puts a subject in perceptual contact with the property (i.e. the experience involves an actual perceptual link with the property). In Gareth’s case, his experience puts him in perceptual contact with the actual shape of his vase. The β -content of his experience would therefore be accurate only if his vase is round.

In order to save the tempting explanation, this gambit must explain perceptual contact without appeal to either perceptual phenomenology or facts about which beliefs it would be rational to form in response to experience. The obvious move is to treat perceptual contact as a purely causal relationship between a subject and an element of the world (an object or property). A subject would be in perceptual contact with a property iff there exists an appropriate causal connection between the subject and the property. For instance, the property may have to be the causal source of information made available to a subject when she undergoes a visual experience.

In making this move, the anticipated theorist already takes a substantive step away from the tempting explanation. For she now recognises a rationally relevant level of perceptual content that is remote from phenomenology. This departure from the tempting explanation weakens what was an attractively tight connection between phenomenology and rationality. I examine the costs of this sort of departure in §5.2.

A deeper problem lurks just around the corner. Suppose we explain β -content in terms of a purely causal notion of perceptual contact. The tolerance of perceptual ascription then becomes easy to accommodate. Gareth’s experience provides him with a wealth of information causally derived from his vase’s actual shape. It therefore puts him in (purely causal) perceptual contact with the

roundness of the vase. Because of the link between this contact and β -content, his experience permits him to rationally believe the vase has *that shape*. So far, so good.

Unfortunately the nice results do not last. Suppose Gareth's round vase were to look triangular rather than oval. His experience would presumably still provide information causally derived from the vase's actual shape. Thus the experience would put him in (purely causal) perceptual contact with the roundness of the vase. Because of the link between this contact and β -content, it permits him to rationally believe the vase has *that shape*, just as it did in the original case.

Wait a moment! We now seem committed to saying that a tolerance case can involve almost any sort of error about a property. But we have already noted that this level of tolerance is implausible: a perceptual link with a property can be misleading while still enabling thought about the property, but it can mislead only so far. I endorsed sortalism about perceptual ascription—the view that to perceptually ascribe a property we must correctly classify it according to its kind—all the way back in §1. Yet sortalism was attractive in part because we recognise non-trivial limits on the tolerance of perceptual ascription.

The ‘causal contact’ account of β -content must therefore impose restrictions on when perceptual contact determines β -content, if it is to respect sortalism. These restrictions will capture non-trivial limits on the tolerance of perceptual ascription. A property will belong in the β -content of an experience iff (1) the experience puts a subject in perceptual contact with the property and (2) the limits of tolerance are not exceeded (in other words, the experience is not *too* misleading). But now the ‘causal contact’ account runs afoul of the explanatory priority that the tempting explanation demands of perceptual content. The limits on tolerance depend upon facts about when a subject can perceptually ascribe a property, and these in turn are facts about what it would be rational to believe in response to an experience. Thus while proponents of the tempting explanation may introduce a

distinction between α -content and β -content, this distinction does not allow perceptual content to perform the role it is assigned by the tempting explanation.

4.2 Deny *Contact and Content*

This subsection discusses the second alternative to my preferred reply to the Tolerance Fork, namely to deny *Contact and Content*. I make good on my promise from §3 to show that a proponent of the tempting explanation who denies *Contact and Content* cannot accommodate the tolerance of perceptual ascription.

Consider Gareth once again. He sees a round vase, but rationally believes it to be oval. He can also rationally form beliefs that perceptually ascribe the actual shape of the vase (those he might express with sentences like ‘it’s that shape’). Being able to rationally form these beliefs in response to his experience is part of what it is for Gareth to be in a tolerance case. However, perceptual ascription requires perceptual contact with the ascribed property: Gareth’s misleading visual experience must therefore put him in perceptual contact with the roundness of his vase.

Now suppose that we deny *Contact and Content*. Gareth retains perceptual contact with the vase’s shape. But this shape no longer figures in the content of his experience of the vase: Gareth’s experience can be accurate even when his vase isn’t round.

The tempting explanation can still explain why Gareth rationally believes that his vase is oval. It explains beliefs rationally formed in response to visual experience by appeal to two features of perceptual content. Gareth’s belief that his vase is oval is *rational* because his experience’s content would be accurate only if the vase were oval. His belief is *false*, despite being rational to form, because his experience’s content is inaccurate.

However the tempting explanation can no longer provide a parallel explanation of Gareth’s beliefs about the actual shape of his vase (the property about which he is misled by his experience). A parallel explanation would require the content of Gareth’s experience to be accurate only if the

vase is round. But the whole point of denying *Contact and Content*—at least for those who want to save the tempting explanation—is to block the move to **2** (the claim that the content of Gareth’s experience is accurate only if his vase is round). Yet the tempting explanation cannot accommodate the tolerance of perceptual ascription if it cannot explain the rational status of Gareth’s beliefs.

5 Abandoning the Tempting Explanation

I’ve considered several replies to the Tolerance Fork on behalf of the tempting explanation. All have run into insuperable difficulties. These difficulties support what I believe to be the best reply to the Tolerance Fork, namely to jettison the tempting explanation.

Once we jettison the tempting explanation, we will need a new explanation of the common sense view (the view that only some beliefs formed in response to experience are rationally formed, and these may be either true or false). The new explanation must avoid the sins of its predecessor. Yet how far must one depart from the tempting explanation to avoid the Tolerance Fork?

The tempting explanation runs afoul of the Tolerance Fork because it entails **4** (the claim that where it is rational to believe that o is F in response to an experience, the content of that experience is accurate only if o is F). The Tolerance Fork arises once we apply **4** to tolerance cases. Applied to Gareth’s case, **4** entails that the content of his experience is accurate only if his vase is oval (i.e. **5**). Yet this result is in tension with **2** (the claim that the content of Gareth’s experience is accurate only if his vase is round).

In what follows, I’ll consider three ways of departing from the tempting explanation that block the move from **4** to **5**: (i) argue that tolerance cases are exceptions to **4**, but insist that **4** otherwise holds; (ii) abandon an element of the tempting explanation yet retain the rest; (iii) abandon the tempting explanation altogether. I shall argue that we ought to pursue (iii) over (i) or (ii).

5.1 Modify the Tempting Explanation?

In order to insist that tolerance cases are exceptions to **4** the tempting explanation must complicate its otherwise straightforward explanation of misleading visual experience. Yet it would be *ad hoc* to simply stipulate that tolerance cases constitute an exception to **4**. Hence we need some independent reason for insisting that tolerance cases are exceptions. Consider the following case:

Dorothy glances at her desk, and sees an oval vase. Yet she also has independent reason for the mistaken belief that (like Gareth) she's seeing the vase through a distorting medium. So in response to seeing the vase, she forms the false belief *that is round*.

As originally developed, the tempting explanation holds that visual experience misleads a subject only if it is illusory. But this claim obviously is not true of Dorothy. A belief about her perceptual situation mediates her response to seeing the vase. As a result of this mediation, she is misled into believing that the vase is round. Yet if her experience possesses content, this content is accurate only if the vase is oval (and so her experience counts as veridical).

In order for the tempting explanation to accommodate cases like Dorothy's it must treat these cases as exceptions to **4**. For when applied to Dorothy's case, **4** entails that the content of her experience is accurate only if her vase is round—an absurd result. What Dorothy's case shows is that a right account of how perceptual experience interfaces with belief must be sensitive to complications arising from the influence of a subject's background epistemic state.²² Whether any particular case counts as an exception to **4** will depend upon the sort of influence exerted by a subject's background epistemic state. In Dorothy's case, she mistakenly believes that her experience is illusory, and this belief imposes a rational constraint on which beliefs she forms in response to her experience (she therefore refuses to believe the vase is oval).

²² Cf. Armstrong (1955); Brewer (2011); Hellie (2011, 2014); Genone (2014). Hellie (2011, pp. 148–9) dubs Dorothy-type cases as 'delusive good' (and later 'suprabad' in §4.1 of Hellie 2014).

Does the existence of cases like Dorothy's provide grounds for thinking that tolerance cases might also be an exception to **4**, and for similar reasons? If it does, the tempting explanation will dodge the Tolerance Fork at the minimal cost of complicating its explanation of misleading visual experience.

Unfortunately the complications Dorothy's case introduces cannot save the tempting explanation from the Tolerance Fork. For Gareth's situation is quite different from Dorothy's.

To begin with, it might not be plausible to insist that Gareth's response to his experience is mediated by background beliefs about his own perceptual situation. The existence of cases like Dorothy's—cases in which mediating beliefs complicate matters—gives us no reason to think that in all cases the path from perception to belief is mediated by other beliefs. In fact, it might be especially implausible to think this in cases like Gareth's, where someone is simply absorbing information delivered by a perceptual link.

Yet even if we shelve these worries about whether other beliefs mediate Gareth's response to his experience, the complications Dorothy's case introduces cannot rescue the tempting explanation from the Tolerance Fork. Insofar as it makes sense to speak of Gareth's (presumably highly tacit) beliefs about his own perceptual situation, his error lies in the other direction: he assumes that he isn't a victim of illusion. Gareth shares this assumption not just with other subjects in tolerance cases, but with all subjects forming *that is F* beliefs by uptake from a perceptual experience whose veridicality they have no reason to doubt. As a result, Gareth's assumption about his perceptual situation cannot make him an exception to **4**. If it could, ordinary instances of belief formation by uptake from perception also count as exceptions. And in that case the tempting explanation melts away as an account of how perception underpins rational belief. So the attempt to minimally modify the tempting explanation in order to avoid the Tolerance Fork threatens to over-generate exceptions to **4**, and so undermine the tempting explanation itself.

I have argued against only one way of treating tolerance cases as exceptions to 4. Ingenious defenders of the tempting explanation therefore remain free to pursue alternative grounds upon which to treat tolerance cases as exceptions. But my discussion brings out a challenge for these determined theorists. On the one hand, their modification of the tempting explanation's explanation of misleading visual experience must not undermine the tempting explanation—a non-trivial obstacle. On the other hand, the modification must also have strong independent support. While some tortured modifications might permit tolerance cases to count as exceptions to 4, such baroque accounts rarely enjoy the kind of deep theoretical or intuitive attraction that made the tempting explanation so tempting in the first place.

5.2 Prune the Tempting Explanation?

As originally presented, the tempting explanation more or less follows from two widely held claims:

- I. What it would be rational to believe in response to a perceptual experience depends upon the conditions under which the experience would be accurate.
- II. The content of a visual experience (its conditions of accuracy) captures how the world looks to a subject undergoing the experience.

(I) captures the core of the tempting explanation.²³ But the tempting explanation can explain the common sense view only if its notion of perceptual content is explanatorily prior to the common

²³ One difference between (I) and (A)—the tempting explanation's claim that whether a belief with a given content counts as a rational response to visual experience depends upon which content the experience possesses—is that the former concerns what it *would* be rational to believe in response to an experience, while the latter concerns whether a belief *is* a rational response to an experience. In order to mirror the distinction between ‘doxastic’ and ‘propositional’ justification (cf. Goldman 1979; Pryor 2000), let’s say that (I) is about ‘propositional rationality’ while (A) is about ‘doxastic rationality’. While this difference would matter under other circumstances, I can ignore it in this chapter. My focus has been beliefs formed in a *standard* fashion in response to information delivered by a perceptual link. The division between ‘propositional’ and ‘doxastic’ rationality matters only once we start to worry about beliefs formed in a *non-standard* fashion. For example, it seems natural to say of Dorothy both that her belief is a rational response to her experience, given the influence of her background beliefs, and that the same belief would not otherwise be a rational response to the experience. One way to capture this bizarre rational status is to say that while Dorothy’s belief is doxastically rational, it is not propositionally rational. Yet my discussion of the relationship between Dorothy’s case and Gareth’s did not require marking the distinction between propositional and doxastic rationality.

sense view. This explanatory priority is secured by (II). Accuracy conditions are introduced to capture an aspect of perceptual phenomenology: two objects can look the same, yet one object may fail to be as it looks. Sameness of accuracy conditions explains the sameness in looks. And two experiences can share their accuracy conditions even when one experience is inaccurate.

We therefore cannot abandon (II) without risking the explanatory priority of perceptual content. Without this explanatory priority, the tempting explanation cannot use (I) as part of an explanation of the common sense view. Yet why not abandon (II) in favour of some alternative non-phenomenological ground for the attribution of perceptual content to visual experience—one that safeguards the explanatory priority of perceptual content over the common sense view?

This gambit quickly runs into trouble. If we were to retain (I) while replacing (II), we would violate an important intuitive connection between perceptual phenomenology and rationality. (Here I pick up a thread from my §4.1.2 discussion of how one might use perceptual contact to construct a second type of perceptual content.) The connection runs roughly as follows. When I undergo a visual experience, the world looks a certain way to me. And the way the world looks to me constrains what it would be rational for me to believe in response to the experience. We lose this connection when we reject (II) while keeping (I). (I) connects rational belief formation to perceptual content, while (II) connects perceptual content to visual phenomenology. Rejecting both of these claims—the option I advocate in the next section—would fully sever the content-based connection between visual phenomenology and rationality, and thus open up room for an alternative connection between the two. In contrast, keeping (I) while replacing (II) with a non-phenomenological ground for the notion of perceptual content would both block the content-based connection, and stymie any attempt to construct an alternative account of the relationship between visual phenomenology and rationality.

5.3 A Radical Outlook

Given the problems raised above, we should pursue radical alternatives to the tempting explanation. These alternatives become available once we reject both (I) and (II).

The Tolerance Fork highlights how the tempting explanation conflates two different ways in which visual experience can ‘fall short’: visual experience can *mislead* us about whether a perceived object possesses a property; it can also *fail to put us in perceptual contact* with an object’s properties (and so prevent us from perceptually ascribing them). Given *Contact and Content*—and ignoring tolerance for a moment—the tempting explanation entails that a visual experience misleads us about whether a perceived object is F just if it fails to put us in perceptual contact with F. (For the view says that misleading experiences are illusory, and content of an illusory experience is accurate only if the perceived object isn’t F.) The tolerance of perceptual ascription forces apart these two ways in which experience can ‘fall short’. It thus exposes the conflation at the core of the tempting explanation.

Instead of trying to salvage some semblance of the tempting explanation, I propose that we pursue radical alternatives that take seriously the separation between the two ways for an experience to ‘fall short’. These alternative accounts must explain the common sense view while respecting the following cross-cutting taxonomy (a taxonomy generated when we ask ‘What would it be rational to believe in response to a given experience?’):²⁴

1 (‘Good’ case)	2 (Tolerance case)	3 (Intolerance case)	4 (Hallucination)
o is G	o is G	o is H	?
Contact with o	Contact with o	Contact with o	No contact with o
Contact with G	Contact with G	No contact with H	?

²⁴ A number of extant accounts may meet these constraints (Cf. Campbell 2002; Hellie 2011, 2014; Kalderon 2011; Brewer 2011; Antony 2011; Johnston 2014; Dickie 2015).

Rational to believe o is G	Rational to believe o is F (where F and G are incompatibles)	Rational to believe o is F (where F and H are incompatibles)	?
Rational to believe o is <i>thus</i> (where ‘thus’ picks out G)	Rational to believe o is <i>thus</i> (where ‘thus’ picks out G)	Not rational to believe o is <i>thus</i> (and ‘thus’ fails to pick out either H or F)	?
Non-Misleading Experience	Misleading Experience		?

Whether an experience counts as misleading or non-misleading—the sixth row of the table—depends both upon how things stand in the world, and upon which beliefs the fourth row says it would be rational to form. At least two kinds of experience count as misleading: ‘tolerance’ cases and ‘intolerance’ cases. (I say ‘at least’ because the taxonomy also includes the mysterious fourth column. This column describes experiences we classify as ‘hallucinatory’. The table highlights their degeneracy: all we can say about hallucinatory experiences at this stage is that they do not provide perceptual contact with o .²⁵)

Experiences that occupy column 3—the ‘intolerance cases’—mislead subjects about which G -type property o has (its shape, for instance), but do not permit these subjects to perceptually ascribe whichever G -type property o actually has (and so do not put subjects in perceptual contact with this property). For example, while Gareth can perceptually ascribe roundness when his round vase looks oval, he would not be able to ascribe triangularity were he to see a triangular vase that looks oval. We must recognise such intolerance cases (and thus non-trivial limits on the tolerance of perceptual ascription) if we’re to accommodate some version of sortalism about perceptual

²⁵ The first row receives a ‘?’ because of the possibility of so-called ‘veridical hallucination’ (cf. Lewis 1980b). The third row has a ‘?’ in order to skirt around an argument offered by Pautz (2010) for (I)—an argument that takes as a premiss the claim that hallucinatory experience permits us to perceptually ascribe properties we hallucinate. The fourth and fifth receive a ‘?’ because I continue to remain neutral about whether we can rationally form beliefs in response to hallucinatory experience. See Martin (2002), Hellie (2010), and Soteriou (2013) for nuanced discussion of how much we can and should say about the nature of hallucinatory experience.

scription. The reader might also have noticed that the fifth row of column 3 says that it would not be rational to believe *that is thus* in response to an experience belonging to that column. This claim does not follow from any earlier argument. The claim simply serves to simplify some of what comes below. However I shall return to this claim in a few paragraphs. It turns out that the significance of my taxonomy remains even if the relevant *that is thus* beliefs count as rationally formed in response to experience.

My taxonomy represents a radical re-fashioning of how we think about visual experience. Because it sharply distinguishes two functions of perceptual experience—to rationalize belief and to put us in perceptual contact with the world—the taxonomy achieves a fine-grained sensitivity that a taxonomy framed in terms of perceptual content cannot match.

We cannot recognise the distinction between the experiences belonging to columns 1, 2, and 3 if we restrict ourselves to the resources made available by the tempting explanation. Proponents of the tempting explanation think there is an underlying notion of perceptual content in terms of which they can explain facts about rational belief formation. We've seen that when they explain this notion of content by appeal to phenomenology—specifically by appeal to the putative possibility of perceptual error—they cannot capture rationally relevant differences between tolerance and intolerance cases (i.e. experiences belonging to columns 2 and 3). Experiences in these cases are indistinguishable from a subject's point of view, which forces proponents of the tempting explanation to assign them the same accuracy conditions. Yet these same experiences stand in different rational relations to beliefs formed in response to them. Of course a proponent of the tempting explanation might instead try to cook up a different notion of perceptual content using just facts about perceptual contact. However it turns out that perceptual contact does no better than perceptual phenomenology when it comes to capturing our taxonomy. For while facts about perceptual contact permit us to distinguish experiences from columns 2 and 3, experiences that belong to column 1 place subjects in perceptual contact with precisely the same entities as

experiences that belong to column 2. As a result, we're left with a taxonomy of perceptual experience that rests on facts about which beliefs it would be rational to form in response to experience.

This ‘rational’ taxonomy forces us to finally throw away the longstanding idea that we can use a distinction between veridical and illusory experience to explain the common sense view. The taxonomy makes available several ways²⁶ one might reconstruct this distinction. But the reconstructed distinction will not be able to perform the explanatory work required by the tempting explanation. For we've seen that no distinction drawn purely at the level of perception can capture facts about which beliefs it would be rational to form in response to experience. More generally, I've shown that we must discard the widespread suggestion that we can use a notion of perceptual content—a notion taken as explanatorily prior to facts about the rational status of beliefs formed by uptake from experience—to explain patterns of rationality for these beliefs. However my argument does not rule out the introduction of *some other* notion of perceptual content (to see how what this might look like, see §1.3 of Chapter 3); the problem lies only with a notion of perceptual content taken as *explanatorily prior* to facts about the rational status of beliefs formed by uptake from experience.

I shall close my discussion by tying up two loose ends. The first concerns how my taxonomy says a subject's belief that his vase has *that shape*—namely triangularity—is not a rational response when the subject's experience belongs to column 3. It turns out that the taxonomy would retain its force even were we to remain neutral on the rational status of this sort of belief. If the subject's

²⁶ Those convinced of the explanatory autonomy of perceptual experience will try to draw the distinction using some other non-content-like feature of visual experience. For example, they might classify illusory visual experience as involving somehow ‘confused’ or ‘unfocused’ perceptual contact with a perceived object (cf. Campbell 2002). By contrast, those less inclined to treat perceptual experience as explanatorily autonomous may draw the distinction in terms of what a visual experience enables a subject to do. For example, perhaps a visual experience is illusory just if beliefs rationally formed in response to the experience would be false, and veridical just if the same beliefs would be true. At their most extreme, these views treat illusory visual experience as a kind of cognitive error akin to a false belief (cf. Armstrong 1955; Austin 1962; Travis 2004; Hellie 2011, 2014; Genone 2014).

belief is a rational response—despite the fact the ‘that shape’ component of his belief fails to pick out his vase’s shape—we lose the rational asymmetry between experiences in columns 2 and 3.²⁷ Instead of a rational asymmetry, the difference rests on a cognitive asymmetry: if the subject undergoes an experience from column 2, he can perceptually ascribe his vase’s shape; if he undergoes an experience from column 3, he cannot. My taxonomy therefore rests upon rational-cum-cognitive facts regardless of the rational status assigned to a subject’s ‘it’s that shape’ beliefs when he undergoes an experience from column 3. As a result, determining the boundaries between columns 2 and 3 will require settling the kinds of questions raised at the end of §1—questions about limits on the tolerance of perceptual ascription. Proponents of the tempting explanation cannot accommodate my rational-cum-cognitive taxonomy any more than my original rational taxonomy. Both taxonomies recognise divisions within visual experience that cannot be captured using perceptual content—or at least not a notion of perceptual content that one could use to explain the common sense view.

The second loose end concerns the tight connection between phenomenology and rationality discussed in §5.2. My taxonomy can respect this connection. For we need only recognise a division between two aspects of the phenomenology of Gareth’s experience: one aspect captures the fact his vase looks oval; the other captures his perceptual contact with the vase’s actual shape. The phenomenology characteristic of perceptual contact with a property is an analogue for properties of the phenomenology characteristic of an attentional ‘lock’ or ‘focus’ on an object. These two aspects of perceptual phenomenology—what we’ll call ‘looks’ and ‘contact’—play different rationality-securing roles. In Gareth’s case, how his vase ‘looks’ makes it rational for him to believe it is oval;

²⁷ This dispute parallels a disagreement in the literature on perceptual demonstrative thought over whether an essayed perceptual demonstrative belief that fails to be about an object can be justified. Campbell (2002) provides an account of perceptual demonstrative thought on which these ‘empty’ beliefs are not justified; Dickie (2015) develops an alternative account on which these beliefs can be justified.

yet the presence of ‘contact’ phenomenology that arises from perceptual contact with the vase’s actual shape helps make it rational for him to believe the oval is *that shape*.

One cannot extract a notion of perceptual content from ‘contact’ phenomenology that can do the explanatory work required by the tempting explanation. Visual experience can possess ‘contact’ phenomenology even when the limits of tolerance are exceeded (just as the phenomenology characteristic of attentional focus may be present even in the absence of an attentional link with an object). Intolerance cases will therefore also sometimes involve ‘contact’ phenomenology. As a result, any notion of content one tries to build out of ‘contact’ phenomenology will fail to carve at the joints with respect to the capacity of an experience to underpin the rational-cognitive facts. The failure of the tempting explanation thus ushers in a new project: to explain how and why these two dimensions of perceptual phenomenology play their rationality-securing roles.

Chapter 2

On The Structure of Predicative Thought

How much of our thought about the world depends upon our being able to both perceive the world and think about what we perceive? One version of this old question concerns thought about objects. It asks how much of our thought about objects (e.g. thought about distant quasars or tiny bosons) depends upon our capacity for perception-based thought about ordinary objects (such as sofas, scarves, and cats). This chapter focusses on a parallel question that concerns thought about properties. It asks how much of our thought about properties (e.g. charge, mass, etc.) depends upon our capacity for perception-based thought about observable properties (e.g. colours, shapes, sizes). Common sense suggests that perception-based thought is more basic. Yet what does this basicness come to? In this chapter I spell out—and argue for—a sense in which perception-based thought about observable properties is more basic than thought about unobservable properties. My argument exploits a deep but neglected connection between thought about properties and our grasp of causation.

Let's begin by reminding ourselves of some terminology already introduced by the last chapter. Suppose a ball bounces into view. This perceptual encounter puts you in a position both to think about the ball and to ascribe bounciness to it, each in a particular way. Call these ways of thinking—of the ball on the one hand, and bounciness on the other—‘perceptual demonstration’ and ‘perceptual ascription’.

The view I shall defend concerns the structure of predicative thought. It mirrors an attractive picture of the structure of thought about objects. This mirrored view treats perceptual demonstration as a form of identification. Let us say that a subject ‘identifies’ an object iff she stands in an identifying relation to it. An identifying relation introduces an object as the object upon whose properties the truth or falsity of (at least one of) a subject’s thoughts depends. Perceptual

demonstration contrasts with another form of identification, which I shall call ‘descriptive identification’. A person identifies an object ‘descriptively’ just if the identifying relation holds in virtue of (a) the subject grasping some descriptive condition and (b) the object being the unique satisfier of this descriptive condition (for example, someone might identify Frege descriptively as the author of the *Begriffsschrift*). This distinction between perceptual demonstration and descriptive identification frames the attractive picture of the structure of thought about objects. On this old but familiar view, our capacity for perceptual demonstration grounds our capacity to descriptively identify objects. Call the view ‘ATOMISM’.¹

Few have investigated whether an analogue of ATOMISM applies to perceptual ascription.² This chapter defends one such analogue:

PREDICATIVE ATOMISM (Approximate Version) Our capacity to ascribe unobservable properties (such as spin or charge) ultimately depends upon our capacity to perceptually ascribe observable properties (such as round or red).

Just as ATOMISM secures an ineliminable role for perceptual demonstration, PREDICATIVE ATOMISM entails that perceptual ascription is essential to our general capacity to ascribe properties (i.e. our general capacity to think thoughts we standardly express with sentences of the form ‘ α is ϕ ’). My argument for PREDICATIVE ATOMISM consolidates the centrality of perceptual ascription. Whereas PREDICATIVE ATOMISM speaks only of thought about properties, my argument also establishes that our grasp of causation partly rests upon our capacity for perceptual ascription of observable properties. The chapter has four parts. §1 clarifies the content and scope of PREDICATIVE ATOMISM.

¹ ATOMISM inherits its name from Russell’s (1918/1956) Logical Atomism. Aside from Russell, defenders of ATOMISM include Strawson (1959), Dummett (1973/1981), Evans (1982), Peacocke (1983), Brewer (1999), Campbell (2009), and Dickie (2010b). As formulated above, ATOMISM remains neutral with respect to the question of whether descriptive identification of an object counts as a way of thinking about an object. This question drives the contemporary debate over whether thoughts involving descriptive identification can count as singular thoughts (Cf. Jeshion 2002 and Recanati 2012).

² Russell (1910; 1918/1956) is one of the few. He suggests that an analogue of ATOMISM applies to the perceptual ascription of properties. Yet Russell’s ability to back up this suggestion has been met with scepticism, most recently in Bostock (2012, Ch. 7).

§2 argues for the view. §3 defends two central premisses of the argument of §2. §4 briefly sketches two additional consequences: the first concerns perception, the second thought. Chapters 3–4 develop and defend these consequences in more detail.

1 Understanding PREDICATIVE ATOMISM

I introduced PREDICATIVE ATOMISM as the claim that our capacity to ascribe unobservable properties (such as spin or charge) ultimately depends upon our capacity to perceptually ascribe observable properties (such as round or red). But this formulation isn't as transparent or precise as the version of PREDICATIVE ATOMISM I plan to defend. This section explains and precisifies the initial formulation. I focus on three aspects of PREDICATIVE ATOMISM: how subjects ascribe unobservable properties; how they perceptually ascribe observable properties; and how the first kind of property ascription depends upon the second.

I shall start by introducing two ways of ascribing properties—what I shall call ascribing a property ‘as dispositional’ and ‘as categorical’. This division roughly parallels a metaphysical distinction between two kinds of property. Metaphysicians divide properties on the basis of what it takes for an object (or some other particular) to instantiate them. A property is ‘dispositional’ when its instantiation conditions are exhausted by an object’s possession of a particular set of causal dispositions or causal powers (e.g. an electron’s tendency to repulse other negatively charged particles). In contrast, a property is ‘categorical’ when its instantiation conditions are not exhausted by these causal dispositions.³

Unlike this metaphysical distinction, our division between two ways of ascribing properties does not depend upon the instantiation conditions of the ascribed properties. Instead it depends upon what a speaker must know in order to ascribe a property. A subject can ascribe a property only

³ For a particularly clear statement of this distinction, see Armstrong (1997, Ch. 5). Many (such as Shoemaker 1980) argue that all properties are dispositional.

if she can identify objects, or other particulars, as bearers of the property (e.g. ‘these chairs are red, those are not’). A kind of classificatory knowledge underwrites this capacity—a sensitivity to conditions under which a property is instantiated. Different ways of ascribing a property require distinct kinds of classificatory knowledge. A subject ascribes a property ‘as dispositional’ iff her classificatory knowledge consists in knowledge of the causal powers conferred by the property.

Both observable and unobservable properties may be ascribed as dispositional.⁴ However unobservable physical properties must be ascribed as dispositional. For we only ever know unobservable properties such as negative charge or spin or polarity through their causal roles. Perceptual experience cannot provide the kind of direct epistemic access that would enable us to ascribe these properties without relying exclusively upon their causal roles. And once we step beyond perceptual experience, we seem left with no option but to identify unobservable properties through their effects on perceptually detectable phenomena. So when classifying something as negatively charged, for example, we can only ever rely upon our knowledge that a negatively charged object will possess certain causal dispositions (e.g. an aversion to positively charged entities).

Furthermore, to ascribe an unobservable property as dispositional requires that a subject be capable of ascribing the causal relations that make up the causal role of the property (notice that these capacities may be interdependent: neither need enjoy priority). The ascription of unobservable properties as dispositional is thus analogous to the descriptive identification of objects in ATOMISM. Just as thinking about an object descriptively requires thinking of the object as the bearer of certain properties, ascribing an unobservable property as dispositional requires ascribing it as the occupant of a certain causal role.

⁴ In the case of observable properties, many of the properties that we see, touch, or hear are ascribed to objects purely on the basis of our knowledge of dispositions these properties confer upon objects that possess them. For instance, Gibson (1977; 1979) argues convincingly that we perceive ways in which objects enable us to act. These so-called ‘affordances’—relations between subjects, an environment, and an object—are dispositional properties. See Siegel (2014) for a recent discussion of challenges raised by affordances for a philosophical account of perceptual experience.

If a subject does not ascribe a property as dispositional, she might instead ascribe it as categorical (the distinction isn't exhaustive). A subject ascribes a property 'as categorical' iff she ascribes it both (1) without relying only upon knowledge of causal powers it confers, and (2) by treating the property as a ground of an object's dispositions (e.g. someone who ascribes circularity as categorical can make sense of a circular object being inclined to roll on smooth inclines). Our primary, if not our only means of acquiring this sort of knowledge of a physical property is perceptual experience. Visual experience, for instance, seems to present some properties directly, unmediated by awareness of causal powers these properties confer upon the objects that instantiate them. Shape and size properties arguably belong in this category. Vehement disagreement persists over whether colours belong to it. Thankfully nothing in this paper hangs on whether we ascribe colours as categorical or as dispositional.⁵ Indeed my definition of what it is to ascribe a property as categorical actually opens up a potential intermediate position. One could deny that we must ascribe colours as dispositional, yet stop short of admitting ascription of them as categorical by insisting that our ascriptions of colour satisfy (1) but not (2).

Metaphysicians may notice that my distinction between ascribing properties as categorical and as dispositional leaves out a way of ascribing properties that parallels the properties posited by Lewis (1986). On Lewis's account, properties are strongly categorical. A 'strongly categorical' property is *categorical* insofar as its identity isn't exhausted by the causal dispositions it confers upon objects that instantiate it. What makes such a property *strongly* categorical is the modal independence of these dispositions and the property's identity. Every subset of the dispositions conferred by a strongly categorical property is such that the property could fail to confer these dispositions. Thus the property of being positively charged, for example, could in some other possible world confer

⁵ However I'm sympathetic to Campbell's (1993, 2006b) defence of the view that we do sometimes perceptually ascribe colours as categorical. Though see Cohen (2010) for an attempt to undermine Campbell's claims about colour phenomenology, and Evans (1985, pp. 268–74) for arguments in favour of the view that we can only ever perceptually ascribe colours as dispositional.

precisely the same causal dispositions as the property of being negatively charged does in the actual world. But what could it be to ascribe a property ‘as strongly categorical’? We cannot appeal, as we did in the case of categorical properties, to how these properties show up in perception. For even if perceptual experience puts us in touch with what are in fact strongly categorical properties, these properties would (like merely categorical properties) show up as categorical grounds of certain dispositions (e.g. round objects roll, heavy objects fall, etc.). Thus perception cannot alone make intelligible the possibility that these perceptually presented properties could occupy a radically different causal role (e.g. as grounds for opposing causal dispositions). Indeed, that attempts to make sense of strongly categorical properties run into this difficulty explains some of the strong intuitive resistance to a Lewisian account of properties. Yet such qualms do not (straightforwardly) undermine Lewis’s posit of strongly categorical properties. To think otherwise would be to confuse a theory of properties with a theory of what it is to ascribe properties in thought.

We can now take a step towards the precise version of PREDICATIVE ATOMISM that I’m going to defend:

PREDICATIVE ATOMISM (Quasi-Precise Version) Our capacity to ascribe unobservable properties depends upon our capacity to perceptually ascribe properties as categorical.

A question remains about the nature of the dependence relation invoked by PREDICATIVE ATOMISM. The relation is not merely nomological. If it were, PREDICATIVE ATOMISM would be an empirical claim about either psychological or developmental priority. While these sorts of empirical claims are of interest, they are not my target. PREDICATIVE ATOMISM is instead a thesis about the necessary structure of understanding. A subject must be able to exercise one capacity—to perceptually ascribe properties as categorical—in order to exercise another—to ascribe unobservable properties.

Yet the dependence relation in PREDICATIVE ATOMISM remains ambiguous. On one construal, any exercise of our capacity to ascribe unobservable properties must involve the exercise

of our capacity to perceptually ascribe properties as categorical. But this version of PREDICATIVE ATOMISM seems implausibly strong. The version that I shall defend is weaker. It says that the dependence relation obtains not between token exercises of the respective capacities, but between each type of exercise. The presence of the one capacity requires the presence of the other, but we can exercise the capacity to ascribe unobservable properties without thereby exercising our capacity to perceptually ascribe properties as categorical. Call this ‘type-dependence’.

We now arrive at our canonical formulation of PREDICATIVE ATOMISM:

PREDICATIVE ATOMISM (Precise Version) Our capacity to ascribe unobservable properties type-depends upon our capacity to perceptually ascribe properties as categorical.

Notice that this thesis carries no commitment to a *metaphysical* asymmetry between observable and unobservable properties. Unobservable properties may still be more fundamental than observable properties.⁶ I thus intend to remain neutral on the question of whether a right metaphysics must explain facts about the one sort of property in terms of facts about the other sort of property.

2 The Argument for PREDICATIVE ATOMISM

This section introduces my argument for PREDICATIVE ATOMISM. I offer an informal statement of the argument, followed by a more careful version. I put off defending two outstanding premisses to the next section.

PREDICATIVE ATOMISM stands at the end of an intuitive line of thought. When we ascribe an unobservable property—such as spin—we exploit our knowledge of its causal role. A property’s causal role is its contribution to the causal behaviour of entities that instantiate it. But to grasp a property as the bearer of such-and-such causal role, we must have a grasp of causation.

⁶ For this notion of fundamentality, see Fine (2001).

We owe this grasp of causation as a generic relation to our grasp of what we'll call 'special causal relations'.⁷ These causal relations involve the ordinary objects we encounter in perceptual experience: a boulder *flattens* a hut, a child *kicks* a ball, a cat *slurps* her water, and so on. We may know of these relations on the basis of experience. It is also difficult, if not impossible, to imagine someone who grasps causation as a generic relation, but who nevertheless cannot ascribe special causal relations (i.e. someone who can think *A caused B*, yet is unable to think anything like *the boulder flattened the hut*). In contrast, nothing strikes us as incoherent about the possibility of a subject who grasps flattening, kicking, or other special causal relations, but cannot ascribe a generic causal relation.

I want to quickly note, if only to dismiss, a misunderstanding that frequently arises at this juncture. Some assume that these claims about our understanding of causation entail the highly contentious thesis that special causal relations are *metaphysically* prior to causal relations that obtain amongst unobservable entities. But this assumption rests on a mistake. As in the case of PREDICATIVE ATOMISM, the priority of special causal relations lies at the level of understanding, not at the level of metaphysics. We need not move directly from this kind of priority to a metaphysical priority thesis.⁸

Returning to the argument: our grasp of special causal relations cannot be separated from either our capacity to think about the ordinary objects they involve, or our capacity to perceptually ascribe observable properties as categorical. Only someone with these two capacities can ascribe special causal relations. PREDICATIVE ATOMISM then follows: our capacity to ascribe unobservable properties type-depends upon our capacity to perceptually ascribe properties as categorical.

Here is a full dress version of the argument:

⁷ Anscombe (1971), later followed by Strawson (1992, Ch. 9), argues for this claim about our grasp of causation.

⁸ Of course, some do think that this account of causal understanding should lead us to endorse a metaphysical priority thesis. Compare Strawson (1992, p. 120).

1. CAUSALISM: Our capacity to ascribe unobservable properties exploits our capacity to ascribe the causal relations that hold in virtue of these properties (i.e. causal relations that constitute the causal roles of these properties).
2. CAUSAL ATOMISM: Our capacity to ascribe causal relations that hold in virtue of unobservable properties type-depends upon our capacity to ascribe special causal relations.
3. Our capacity to ascribe unobservable properties type-depends upon our capacity to ascribe special causal relations. [From 1 and 2]
4. CONJECTURE: Our capacity to ascribe special causal relations type-depends upon our capacity to perceptually demonstrate objects (i.e. to identify objects solely on the basis of perceptual encounters with them).⁹
5. Our capacity to ascribe unobservable properties type-depends upon our capacity to perceptually demonstrate objects. [From 3 and 4]

But what might our capacity to ascribe special causal relations require *beyond* a capacity to identify objects through perceptual demonstration? In part by relying upon 5, I'll argue for the following answer, one that bridges the gap between 1–5 and PREDICATIVE ATOMISM:

6. BRIDGE PREMISS: Our capacity to ascribe special causal relations also type-depends upon our capacity to perceptually ascribe observable properties as categorical.¹⁰

3 and 6 entail PREDICATIVE ATOMISM:

7. PREDICATIVE ATOMISM: our capacity to ascribe unobservable properties type-depends upon our capacity to perceptually ascribe observable properties as categorical. [From 3 and 6]

⁹ Cf. Campbell (2002, p. 249).

¹⁰ Others have defended something like the BRIDGE PREMISS. Strawson (1992, p. 122) appears to defend the nearby claim that our capacity to ascribe special causal relations type-depends upon our capacity to perceptually ascribe observable properties as dispositional (though Roessler 2011 argues that Strawson should have appealed instead to our capacity to ascribe these properties as categorical). Also, Peacocke (2011, p. 162–3) endorses something like the BRIDGE PREMISS (though he doesn't argue for it).

In upholding the entailment from 3 and 6 to 7, I'm assuming that type-dependence is transitive. This transitivity should not be controversial. For I introduced type-dependence to capture what should be a transitive explanatory connection between cognitive capacities. However not every dependence relation in the neighbourhood of type-dependence is transitive (causation, for instance, may not be transitive).¹¹

Before I go on to defend this argument in detail—the task of the next section—I wish to briefly register an observation about the relationship between my argument and an old debate (going back at least to Hume 1739) over whether we perceive causation. Hume argued that perception does not present causal relations. To think otherwise, Hume claimed, is to mistake perception of mere regularity for perception of genuine causality. But a growing cadre has sought to resist Hume's insistence that we do not perceive causation.¹² My argument for PREDICATIVE ATOMISM takes no side in this debate. For PREDICATIVE ATOMISM is a claim about the structure of thought, not a claim about what we can perceive. My argument for the thesis turns on claims about the requirements on thinking certain kinds of thoughts. But whether we perceive causation or not fails to tell us if our capacity to ascribe special causal relations depends upon our capacity to perceptually ascribe properties as categorical. To see or feel causation may be one thing, a capacity to ascribe special causal relations—for instance to think of one object as flattening (or denting, or pushing...) a second object—may be quite another. Furthermore, special causal relations are not distinguished by being perceived, but by the fact that subjects may know about them on the basis of perception.

¹¹ Lewis (1973) builds his counterfactual account of causation around the assumption that causation is transitive, which is why he does not identify causation with the non-transitive relation of causal dependence. But as Lewis (2000, §IX) later came to recognise, counterexamples put serious pressure on this assumption of transitivity.

¹² These rebels include Anscombe (1971), Strawson (1992, Ch. 9), and Siegel (2010).

3 A Foundation for PREDICATIVE ATOMISM

I've already briefly defended two of the four premisses that support PREDICATIVE ATOMISM: CAUSALISM in §1 when I first introduced the ascription of properties as dispositional, and CAUSAL ATOMISM at the beginning of §2. In this section I defend the remaining two premisses: CONJECTURE and the BRIDGE PREMISS. More could be said in defence of the other two premisses. I focus on CONJECTURE and the BRIDGE PREMISS because CONJECTURE is a more controversial and unfamiliar claim, and because the BRIDGE PREMISS has yet to receive a decisive defence. My defence of the BRIDGE PREMISS is my principal positive contribution to the argument for PREDICATIVE ATOMISM. It is also this aspect of my argument for PREDICATIVE ATOMISM that generates many of the implications for perception and thought.

3.1 CONJECTURE

CONJECTURE is the claim that our capacity to ascribe special causal relations type-depends upon our capacity to perceptually demonstrate objects. The best argument for this claim lies in the fact that it answers an old puzzle, introduced by Shoemaker (1988, pp. 211-15), which challenges any view in which our capacity to think about objects rests on our capacity to ascribe causal relations.

Shoemaker's begins with a question: what resources are required to capture our ordinary conception of physical objects (e.g. tables, planets, etc.)? We conceive of the class of physical objects as a non-arbitrary ontological category. Thus a chair counts as a physical object, as does a boulder, but the mereological sum of the chair and the boulder does not. But if we start with a plenitudinous picture in which we think of objects as defined geometrically in terms of totalities of actual and possible space-time trajectories—areas of space whose contours change over time—we cannot generate a non-arbitrary reason to think that the objects picked out by our ordinary conception are

special.¹³ Consider the set of space-time points that specifies an object's trajectory in the actual world. There exists a set of points just like the original, but which begins five minutes earlier. Call this collection of points, which counts as an object under the plenitudinous view, a 'shadow' of the original object. We can define counterparts of a shadow object in every possible world in which the original object exists. Furthermore, one could repeat this shadow-making process for every object with which we interact, and for every object that these objects interact with, and so on. On a plenitudinous picture of objects, we have no non-arbitrary grounds for excluding these shadows from the class of physical objects.

The obvious next step, after a plenitudinous view, would be to claim that physical objects have a particular kind of causal unity. It is thinking of objects in terms of their causal role that supposedly supports a conception of physical object that allows us to non-arbitrarily distinguish physical objects from their 'shadows'. Yet this suggestion also quickly runs into difficulty. The shadow objects will exhibit the same regularities in the actual world, and their counterparts the same regularities across possible worlds, as the ordinary objects and their counterparts. So if a regularity or counterfactual account of causation were true, to think about objects as causal unities would fail to privilege physical objects over their shadows.¹⁴

¹³ This plenitudinous view has its roots in Quine (1960, §36).

¹⁴ To see why this holds, consider Lewis's (1973) counterfactual account of causation. His account has two parts. First, an event e causally depends on an event c iff (1) had c not occurred, e wouldn't have; and (2) had c occurred, e would have. Second, c causes e iff there is a finite chain of events $e_1 \dots e_n$ that is such that e_1 causally depends on c , e_2 causally depends on e_1 , ... and so on until e causally depends on e_n . Now let o be a physical object, and o_s its shadow. If E is the set of events involving o (e.g. o 's being pushed into another object o_1), then there exists a corresponding class of shadow events E_s involving o_s (i.e. if an event involving o also includes another object o_1 , the event's shadow will involve o_1 's shadow; and if the first event starts at time t , the second will start at a 'shadow' time t_s ; and so on). Any counterfactual true of $e \in E$ will also be true of a corresponding event $e_s \in E_s$, since the counterfactuals true of events in E depend for their truth on the behaviour of counterparts of o (and other physical objects), and this behaviour will be mirrored by counterparts of o_s (and the other shadow objects). It follows, on Lewis's account, that if $e \in E$, and e causally depends on c , then there will be an event $e_s \in E_s$ that causally depends on c_s (the shadow of c). And so if $e \in E$, and e causes c (or c causes e), then there will exist an event $e_s \in E_s$ such that e_s causes c_s (or c_s causes e_s). Thus any causal unity displayed by an ordinary object will be matched, on Lewis's account of causation, by a parallel causal unity displayed by the object's shadow.

Campbell (2002, p. 247) argues that Shoemaker's puzzle arises because we assume that the question of which objects exist must be settled by appeal to causal relations we can ascribe without thinking about objects (except as the occupants of certain causal roles). He proposes we drop this assumption. In its place, we should instead hold that our capacity to ascribe causal relations type-depends upon our capacity to perceptually demonstrate objects. Perceptual demonstration exploits our perceptual encounters with objects, rather than our grasp of a descriptive condition. It therefore enables us to identify objects without merely identifying them as the occupants of certain causal roles. CONJECTURE then follows (given CAUSAL ATOMISM): our capacity to ascribe special causal relations type-depends upon our capacity to perceptually demonstrate objects. And given CONJECTURE, our capacity to perceptually demonstrate objects is part of what enables us to recognise objects as causally relevant.

One way to block CONJECTURE would be to offer an alternative to Campbell's solution to Shoemaker's puzzle. We needn't look far in search of a menu of potential solutions. The puzzle bears structural similarities to Putnam's model-theoretic arguments against realism.¹⁵ Putnam exploits the fact that a language construed as a consistent set of uninterpreted sentences will always have more than one model (i.e. a semantic interpretation that assigns semantic values to expressions in a language such that every sentence in the language come out true). Similarly, our puzzle relies upon the possibility of multiple models for our talk of physical objects and causation. All of our causal/objectual claims come out true both on a physical object model and on a shadow object model. Viewed through the lens of Putnam's model-theoretic argument, Campbell's response to our puzzle is to insist that our capacity to think about physical objects rests on our grasp of the objects themselves. Our capacity to think about an object is not the product of an interpretation. It instead

¹⁵ Cf. Putnam (1981). For a recent elaboration and defence of Putnam's arguments, see Button (2013).

occurs prior to any process of interpretation. Our thinking of a particular object becomes part of the input to interpretation, not a result our interpretation must secure.

In light of this connection to Putnam, those seeking an alternative solution to Shoemaker's puzzle might appeal to a version of Lewis's (1983) notion of a 'natural property'. Lewis introduces natural properties to play a range of property-theoretic roles: in explaining causation, in explaining supervenience, and so on. One might well think that an appeal to natural properties will solve the puzzle. The flagship claim of any such solution must be that the property of being a particular ordinary object—a chair for instance—always counts as more natural than the property of being a corresponding 'shadow' object. Thus an appeal to a capacity to ascribe natural properties solves our puzzle only if it enables us to classify 'shadow objects' as less natural than ordinary objects. But would a shadow object really count as less natural than the corresponding ordinary object? Or, staying within the domain of properties, would the property of being a shadow object really count as less natural than that of being an ordinary object? On the account of gradable naturalness hinted at by Lewis (later adopted or adapted by others) the naturalness of a property is a function of the complexity needed to define a predicate that expresses the property using only predicates that express 'perfectly' natural properties.¹⁶ On this account, a property Φ is more natural than a property Ψ just in case some predicate that expresses Φ can be defined more simply, using only predicates that express perfectly natural properties, than any predicate expressing Ψ .

Unfortunately this sort of solution to Shoemaker's puzzle runs into trouble even if we accept all this Lewisian machinery. Any notion of naturalness which makes naturalness a matter of structural complexity is unlikely to work here because of the isomorphism between shadow objects and ordinary objects. For it just isn't clear that a graded notion of naturalness will assign shadow objects and ordinary objects (or rather the properties of being each of these things) different degrees

¹⁶ Cf. Lewis (1984, p. 228). See Williams (2007, §2) for a critical discussion of graded naturalness.

of naturalness. Shadow objects and ordinary objects are isomorphic with respect to everything but their temporal profile. And surely the time when an object comes into being (or when it ceases to exist) cannot affect the complexity of the definition of a predicate that expresses the property of being that object. Thus an appeal to a capacity to ascribe natural properties as dispositional will not alone solve the puzzle, unless we move to a radically different criterion for relative naturalness.

The Lewisian appeal to natural properties does not exhaust the potential alternatives to Campbell's solution. Putnam's opponents have developed a menu of potential replies to his model-theoretic arguments against realism. Doubtless a dedicated theorist could construct analogues of these replies as solutions to Shoemaker's puzzle. But I'm confident that Campbell has nailed down the right solution. So I shall say no more in defence of CONJECTURE. My primary contribution to the case for PREDICATIVE ATOMISM will be the argument I develop for the BRIDGE PREMISS in the next section.

3.2 BRIDGE PREMISS

The BRIDGE PREMISS is the claim that our capacity to ascribe special causal relations type-depends upon our capacity to perceptually ascribe observable properties as categorical. I shall defend it by arguing for three nested component claims: (1) our capacity to ascribe special causal relations type-depends upon our capacity to ascribe observable properties; (2) this ascription must be perceptual; and (3) we must ascribe at least some of the properties as categorical. I shall defend each claim in turn.

3.2.1 Our capacity to ascribe special causal relations type-depends upon our capacity to ascribe properties

One way to argue for (1)—the claim that our capacity to ascribe special causal relations type-depends upon our capacity to ascribe properties—would be to first assume a particular conception of causation. We could then argue that the ascription of special causal relations, construed according

to our assumed conception of causation, requires that a thinker have the capacity to ascribe properties. To ascribe a relation is to think that an entity stands in the relation to some other entity (e.g. that the table is next to the chair). And plausibly a subject can think this sort of thought only if she can think about (or ascribe) the entities the ascribed relation relates. Consequently, (1) will be nearly trivial if we assume a conception of causation on which causal relations relate properties (rather than events, states of affairs, or objects).

As trivial as (1) might seem on this account of causation, it appears much less plausible on others. For instance, what if causal relations relate events (e.g. the boulder falling) rather than constituents of these events (i.e. objects and properties)? Our capacity to think about an event may not depend upon our capacity to ascribe properties (e.g. it may be possible to perceptually demonstrate events). So it appears that one could accept that ascription of a causal relation requires a capacity to think about its relata, yet at the same time deny (1).

However (1) would be considerably less interesting if it were found to rest on a particular account of what causal relations relate. This dependence would also undermine the interest of PREDICATIVE ATOMISM. My goal in this subsection is to argue that acceptance of (1) does not depend upon what we take causal relations to relate. I show that it has a deeper anchor.

An argument for (1)—the claim that our capacity to ascribe special causal relations type-depends upon our capacity to ascribe properties—that does not depend upon a particular account of what causal relations relate must spell out a connection between our capacity to ascribe special causal relations, on the one hand, and our capacity to perceptually ascribe properties, on the other. This connection needs to be quite general: it must hold of both our sophisticated ascriptions of causal relations and our most primitive.

My argument for (1) rests on a minimal constraint on our general capacity to ascribe relations in thought. I contend that the link between property ascription and the ascription of special causal relations has its roots in the generality implicit in any ascription of a relation. Consider the

Generality Constraint defended by Evans (1982): necessarily, to count as thinking that a is F , a thinker must also have the capacities to think that a is H , a is G , and so on (for some suitable range of H , G , etc.), and to think that b is F , c is F , and so on (for some suitable range b , c , etc.). This constraint extends to the ascription of relations: necessarily, to count as thinking that a stands in R to b , a thinker must also have the capacities to think that a stands in R_2 to b , a stands in R_3 to b , and so on (for some suitable range of R_2 , R_3 , etc.), and to think that a_1 stands in R to b_1 , a_2 stands in R to b_2 , and so on (for some suitable range of pairs a_1 and b_1 , a_2 and b_2 , etc.).

Controversy persists over whether and how to restrict the two ranges of thoughts.¹⁷ But my argument relies on only a minimal constraint entailed by any relational version of Evans's Generality Constraint. This minimal constraint is the claim that a subject counts as ascribing a relation only if she could apply this relation to other relata. Put formally: necessarily, to think that a stands in R to b , a subject must have the capacity to think that a_1 stands in R to b_1 (for at least one pair a_1 and b_1 distinct from the pair of a and b). Applied to special causal relations, my minimal constraint entails that to ascribe a particular special causal relation a subject must know what it would be for a different causal interaction to involve that same causal relation. But to apply a relation to other causal interactions in this fashion a subject must grasp the relation as *repeatable*.

What does it take to grasp a special causal relation as repeatable? Given CONJECTURE, to ascribe a special causal relation requires a capacity to perceptually demonstrate objects, and so an ability to recognise these objects as causally relevant. But objects enter into a wide range of special causal relations: an object can flatten, disperse, discolour, and so on. So to ascribe a particular special causal relation as repeatable a subject must be able to know what makes it the case that *this* relation,

¹⁷ Recent advocates of (versions of) the Generality Constraint include Campbell (1986), Peacocke (1992), Camp (2004), Dickie (2010a), Beck (2012), and Recanati (2012). Evans restricts the ranges of thoughts to those that obey categorial restrictions (e.g. we need not be able to think that Julius Caesar is prime), a restriction Camp (2004) resists. Dickie (2010a) pushes in the other direction: she shows that extant arguments fail to establish even Evans's restricted constraint.

rather than some other one, obtains from one moment to the next (e.g. what makes it the case that, during a given interval of time, an object is flattening a hut rather than toppling it).

Whether one special causal relation or another obtains over an interval depends on the sorts of changes involved. Some relations hold in virtue of one set of properties, others in virtue of different properties. Change some of these properties, and one sort of causal relation will obtain; change another set of properties, and a different sort of causal relation may result. For example, a causal interaction counts as an instance of flattening in virtue of properties of both the flattening object (the *agent*) and the flattened object (the *patient*). A subset of the agent's properties brings about a particular change in the shape and size of the patient (i.e. the patient is flattened). If a different subset of properties of the patient were changed, the causal interaction wouldn't be a flattening (it might instead be a discolouring, if the patient's colours change).

I'm now in a position to offer my argument for the claim that our capacity to ascribe special causal relations type-depends upon our capacity to ascribe properties. We begin with the application of my minimal constraint to the ascription of special causal relations:

1. For a subject to ascribe a particular special causal relation she must know what it would be for a different causal interaction to involve that same causal relation—she must grasp the relation as *repeatable*.
2. But to grasp a special causal relation as repeatable requires that a subject be able to think, of some other causal interaction, that they involve instances of the ascribed causal relation (for instance, to grasp what it is for one thing to *push* another, I must be able to think, of some other scenario, that it involves one thing pushing another).

However

3. Being able to think this way requires that a subject distinguish (in thought) that relation from other special causal relations.

After all, I cannot entertain the thought *Judy pushes Jeremy* unless I can distinguish (in thought) Judy pushing Jeremy from Judy tripping Jeremy. But

4. To isolate a *particular* special causal relation a subject must distinguish a *particular* dimension of change.

1–4 generate an initial conclusion about the ascription of special causal relations:

5. In order to ascribe a particular special causal relation a subject must be able to distinguish a *particular* dimension of change. [From 1–4]

Yet

6. For a subject to distinguish a particular dimension of change it is not sufficient to think of an agent or patient as *somewhat* altered; she must think of the agent or patient as changed in a particular way.

Notice that 6—and thus my argument as a whole—does not presuppose a particular conception of what special causal relations relate. While the changes must involve the objects that serve as agent and patient, 1–6 remain silent as to the relationship between these objects and the relata of special causal relations. The relata could be events that have these objects as parts, or facts with the objects as constituents; they could even be the objects themselves. So my argument preserves an important degree of neutrality about the nature of causation. However we still require a gloss on what it is for an agent or patient to change in a particular way:

7. These ways for objects to change are distinguished by the properties these changes involve.

If we want to think of an agent or patient as changed in a particular way, we must be able to think about what distinguishes this way of being changed.

So 6 and 7 entail 8:

8. To distinguish a particular dimension of change a subject must be able to ascribe properties.

[From 6 and 7]

And 5 and 8 entail 9:

9. A subject who ascribes special causal relations must be able to ascribe properties. [From 5 and 8]

If ϕ -ing requires being able to X , our capacity to ϕ type-depends on our capacity to X . So 9 entails (1): the claim that our capacity to ascribe special causal relations type-depends upon our capacity to ascribe properties.

This argument represents a real step forward.¹⁸ It establishes a deep connection between the ascription of properties and the ascription of special causal relations—a connection that remains of independent interest even when we ignore its link to PREDICATIVE ATOMISM. The argument also avoids presupposing a specific account of what special causal relations relate. It instead derives the connection from a general claim about the conditions under which a subject can ascribe a relation.

3.2.2 Our capacity to ascribe special causal relations type-depends upon our capacity to perceptually ascribe properties

I've shown that our capacity to ascribe special causal relations type-depends upon our capacity to ascribe properties. The current subsection takes us from this result to the claim that our capacity to ascribe special causal relations type-depends upon our capacity to *perceptually* ascribe properties.

Given the argument of the previous subsection, we know our capacity to ascribe special causal relations depends upon our capacity to ascribe *some* properties. What might these properties be? 1–5 already entails that they cannot be unobservable physical properties. We can only ascribe unobservable physical properties as dispositional. And our capacity to ascribe these properties as dispositional rests upon our capacity to ascribe special causal relations. Hence the latter capacity cannot depend upon the former.

¹⁸ My argument improves upon attempts to forge the requisite connection by appeal to the role of counterfactuals in ascriptions of special causal relations. Roessler (2011, pp. 84–5) appeals to this kind of link between causal ascription and counterfactuals to motivate a version of (1) on behalf of Strawson (1992) (for discussion of the connection between our grasp of counterfactuals and our grasp of causation, see Hoerl et al. 2011).

The ascription of special causal relations also cannot rest on our ascription of non-physical properties (e.g. properties of abstract objects, like being a number, or properties of mental states, like the property of being rational). We might be able to ascribe these properties without perceiving them. But our ascription of them cannot figure in grasp of special causal relations as repeatable. Non-physical properties like being a number or being rational standardly do not confer causal powers upon objects that instantiate them. As a result, we cannot use them to isolate the dimensions of change characteristic of a specific special causal relation.

We must therefore look elsewhere for a class of properties whose ascription underwrites our grasp of special causal relations. The class of properties in question is not the class of unobservable physical properties. It is also not the class of non-physical properties. Our only alternative is to admit that grasp of special causal relations rests upon our capacity to ascribe observable properties. What distinguishes the ascription of observable and unobservable properties is precisely that a subject must, at least in the first instance, perceptually ascribe observable properties. Thus the argument from 1–5 provides direct support for the claim that our capacity to ascribe special causal relations type-depends upon our capacity to perceptually ascribe observable properties.

Another argument for this claim becomes available if the line of thought developed in the next subsection goes through. There I shall argue that we must perceptually ascribe properties as categorical if we're to ascribe special causal relations. Once that claim is in hand, to falsify the BRIDGE PREMISS a subject would have to ascribe a physical property as categorical without relying upon either perception of the property or a capacity to perceptually ascribe other properties as categorical. Yet someone can ascribe a property as categorical only if she possesses a suitable kind of classificatory knowledge: her sensitivity to the instantiation of the property cannot be exhausted by her knowledge of its causal role. This much follows from our original definition of what it is to ascribe a property as categorical. But from whence could such knowledge derive, if not perceptual experience of the properties themselves?

A temptation will strike some opponents at this juncture. They will appeal to examples usually assumed to undermine broadly empiricist accounts of concept possession, such as Hume's missing shade of blue. Hume imagines a subject who has seen all but one of the visible shades of blue. He claims that if these shades were placed before the subject in order of blueness, the subject would be able to think about the missing shade of blue (despite never having seen it).¹⁹ But while such examples undermine positions on which 'more complex' property ascriptions must be analysed in terms of perceptual ascription, they do not undermine a view on which the ascription of properties as categorical type-depends upon perceptual ascription of properties as categorical. A subject can ascribe Hume's missing shade of blue, for example, because of a prior capacity to perceptually ascribe a range of colours. Similar observations hold of potential parallels to the missing shade of blue for shape, size, and other observable properties. So we can accept that subjects may ascribe some observable properties as categorical in the absence of perceptual encounters with these properties, yet deny that these ascriptions are more basic than perceptual ascription of properties as categorical.

3.2.3 Our capacity to ascribe special causal relations type-depends upon our capacity to ascribe properties as categorical

This subsection lays the final stone on the road to the BRIDGE PREMISS. We must show that the ascription of special causal relations type-depends on the ascription of properties as categorical. I offer two arguments: the first relies upon a claim about our cognitive limitations; the second reaches for a deeper anchor.

Argument From Cognitive Limitations Consider what it would take for the target claim to be false. Our capacity to ascribe special causal relations would merely require a combination of our

¹⁹ Cf. Hume (1739, 1.1.1.10).

capacity to ascribe properties as dispositional and our capacity to perceptually demonstrate objects. As we've already seen, our capacity to ascribe special causal relations depends upon our capacity to ascribe properties, in part because the latter capacity enables us to isolate *particular* special causal relations. I shall argue that this role for property ascription cannot be fulfilled by a capacity to ascribe properties as dispositional.

A special causal relation is the realisation of a causal power possessed by an object. If a boulder flattens a hut, for example, this flattening is the realisation of a causal power belonging to the boulder: roughly, to flatten objects of such-and-such a size, when dropped from a certain height, when the flattened object is made of suitably fragile material, etc. As this example makes vivid, any proper specification of the causal power must be enormously complicated, since whether the causal relation obtains will be conditional upon myriad complex background conditions.²⁰

Now for our ascription of properties as dispositional to underwrite our capacity to distinguish particular special causal relations, it would have to be within our cognitive reach to isolate whichever highly conditional causal powers distinguish a particular special causal relation from other such relations. To isolate these causal powers would require knowledge of how their realisation is sensitive to changes in background conditions. And this does not seem like something a creature with our limited cognitive abilities could do. We can ascribe properties as dispositional without possessing the kind of detailed conditional knowledge required to isolate one of the causal powers conferred by an ascribed property. But without this fine-grained knowledge these property ascriptions will not underwrite any ascription of a special causal relation.

This issue of cognitive load does not afflict the ascription of properties as categorical. The problem with ascribing properties as dispositional was that we could not isolate the specific

²⁰ On the account of dispositional properties developed by Shoemaker (1980), an object has a power conditional on a property or set of properties Q if it has some other property Φ such that the combination of Q and Φ is causally sufficient for having the power, while Q is not alone causally sufficient for having the power. Shoemaker's account treats dispositional properties as clusters of these conditional causal powers.

contribution of a property to change (and thus causation) without detailed knowledge of the conditions that must be met for the associated causal powers to be realised. In contrast, isolating the contribution of a property ascribed as categorical does not require that we inquire into the detailed realisation conditions of any associated causal powers. To consider a property as categorical is to consider it as something which acts as an autonomous basis of change and causation. For example, when a wheel rolls down a hill, an adequate account of what caused this behaviour may simply appeal to the roundness of the wheel. So only the ascription of properties as categorical can underwrite our ascription of special causal relations—at least for subjects with limited cognitive resources.

An Alternative Argument While my first argument gets us the conclusion we want, I suspect we can do better. In what follows I develop an alternative argument that eschews reliance upon claims about our limited cognitive resources. If successful, this argument establishes an even more secure foundation for the target claim, namely the claim that ascription of a special causal relation requires ascription of a property as categorical.

My first argument exploited the fact that any adequate specification of a causal power must be enormously complicated (since whether the causal relation obtains will be conditional upon myriad complex background conditions). Following Shoemaker (1980), call these background-sensitive causal powers ‘conditional causal powers’. The inevitable complexity of a proper specification of a conditional causal power prevents agents with our limited cognitive resources from acquiring sufficiently detailed knowledge of the cluster of causal powers conferred by a property.

I shall argue that this complexity introduces yet another obstacle for subjects who ascribe properties as dispositional—an obstacle that would face even those without our limited cognitive resources. Suppose I come to think that a rolling ball pushed a cup. To think that one object pushes another is to ascribe a special causal relation—pushing. We’ve already seen that to ascribe this

relation I must know what it would be for some other pair of objects to be such that one pushes the other. Yet knowledge requires non-luckiness: when thinking of some other situation involving a pair of objects, my ability to isolate their interaction as an instance of pushing (rather than some other relation) cannot rest on mere luck. For instance, if I think of a pair of objects interacting in a certain way—where this interaction actually constitutes an instance of *pulling* rather than *pushing*—I must count as unlucky if I think of it as an instance of pushing. (To be clear: the non-luckiness involved concerns a subject’s ability to isolate a relation *in thought*, not her ability to recognise instances of this relation out in the world on the basis of perception, testimony, or the like.)

Yet what does it take for a subject to non-luckily isolate a special causal relation? Since this capacity rests on our ability to isolate (again in thought) the changes characteristic of the ascribed special causal relation, our isolation of these changes must also count as non-lucky. But we isolate these changes in terms of the properties they involve. So isolating them in thought requires an ability to non-luckily isolate the relevant properties—and (given our story about the ascription of special causal relations) to do so within a range of potential causal interactions. For example, suppose that the change characteristic of pushing involves a certain motion property. Isolating the former in thought will require isolating the latter. And this isolation will require non-luckily distinguishing that motion property from similar (but distinct) motion properties whose presence in a potential causal interaction would not help constitute the sort of change characteristic of pushing.

We now have a new constraint on the kind of property ascription that must underwrite our capacity to ascribe special causal relations: the classificatory knowledge characteristic of this kind of property ascription must enable a subject to non-luckily isolate a property in thought within a diverse range of potential causal interactions. In particular, the classificatory knowledge cannot leave open rationally-relevant potential causal interactions in which the property a subject isolates is not the ascribed property. These potential causal interactions will often involve quite diverse background conditions. As a result, the conditional causal powers manifested in these cases will vary widely.

I shall argue that only our capacity to ascribe properties as categorical can satisfy the new constraint. Hence only this capacity can underwrite our capacity to ascribe special causal relations. To begin with:

1. Suppose all we have is a capacity to ascribe properties as dispositional.

It follows that

2. We isolate properties in thought by treating them as clusters of conditional causal powers.

[From 1]

But we've seen that

3. Being able to ascribe a special causal relation requires being able to non-luckily isolate properties across a diverse range of potential causal interactions.

Now assume for the moment that

4. When attempting to isolate a property F within a potential causal interaction using a subset of its conditional causal powers, there will always be rationally-relevant potential causal interactions in which a non-F property manifests that subset of F's conditional causal powers.

A property manifests only a subset of its associated conditional causal powers in any given situation (we cannot after all have every condition operative at once). But then given that our classificatory knowledge exploits only a subset of F's conditional causal powers, 2 and 4 entail that attempts to use this classificatory knowledge to isolate F in thought will fail to eliminate rationally-relevant potential causal interactions which involve non-F properties that manifest the relevant subset of F's conditional causal powers. So:

5. If a subject attempts to isolate F within a diverse range of potential causal interactions, she will fail to non-luckily isolate F.

If ascribing a property as dispositional prevents a subject from being able to non-luckily isolate other instances of the property across a range of potential causal interactions, 3 entails that this kind of property ascription cannot underwrite our capacity to ascribe special causal relations.

Note that 5 does not prevent subjects from ascribing properties as dispositional. For the range of potential circumstances within which we must isolate a property when simply ascribing it as dispositional—as opposed to ascribing it for the purposes of ascribing special causal relations—need not involve the kind of variation in background conditions that at 5 undermines non-luckiness.²¹ All I've shown is that our capacity to ascribe properties as dispositional cannot underwrite our ascription of special causal relations. Nothing prevents this capacity from performing other explanatory roles.

In contrast to those who must ascribe properties as dispositional, a subject who can also ascribe properties as categorical faces no parallel obstacle to their capacity to ascribe special causal relations. If a subject ascribes a property as categorical, isolating this property in thought will rarely depend on luck. The classificatory knowledge that underwrites the ascription of a property as categorical enables a subject to non-luckily isolate the property within a diverse range of potential causal interactions. For this knowledge isn't restricted to causal powers a property might share (in whole or in part) with others. As a result, a subject who ascribes a property as categorical can non-luckily isolate the property within a range of potential causal interactions—an ability that will enable her to ascribe special causal relations.

This second argument has been more involved than the first. Yet I believe it succeeds in strengthening my case for the claim that ascription of special causal relations requires a capacity to

²¹ Though see Kalderon (2008) for a complex discussion of the relationship between how colours look (the manifestation of one of a colour's conditional causal powers) and our knowledge of which colours we see. He emphasises the dynamic character of knowledge-which: how a colour looks at a given time may not settle which colour it is; we must often wait for it to manifest other looks. Hence successful perceptual ascription of a colour as dispositional may require a sustained perceptual interaction with the colour—one that reveals a non-trivial subset of the conditional causal powers associated with the colour.

ascribe properties as categorical. Apart from its role in securing PREDICATIVE ATOMISM, this claim will play a major role in Chapter 3. For it turns out that many accounts of perceptual experience cannot accommodate perceptual ascription of observable properties as categorical.

An Assumption Vindicated A loose end remains from my second argument: we require some reason to accept 4. This was the assumption that when attempting to isolate a property F within a potential causal interaction using a subset of its conditional causal powers, there will always be rationally-relevant potential causal interactions in which a non-F property manifests that subset of F's conditional causal powers.

Suppose I want to use a subset of a property's conditional causal powers to isolate it within a diverse range of potential causal interactions. The members of this subset will also belong to other clusters of conditional causal powers—each of which individuates a property. Some of these properties will be quite bizarre and gerrymandered (think grue-type²² properties). Causal interactions involving them will almost always count as rationally irrelevant for subjects who seek to isolate the original property within a diverse range of potential causal interactions.

However non-bizarre properties can also share the subset of conditional causal powers. For instance, nearby determinates of a determinable property will share a many of the original property's conditional causal powers. If these determinates are reasonably different (e.g. circularity and marginal-oval-ness), subjects will often be able to discriminate them by the conditional causal powers they manifest. As a result, these nearby determinates will count as rationally relevant. Unlike the bizarre properties, they cannot be ignored on the grounds of being gerrymandered. And unlike rarely instantiated properties that share causal powers with an ascribed property—and which are rationally irrelevant due to their rarity—these properties are pervasive. Hence they belong to the

²² Grue-type properties depend upon time in a seemingly artificial or arbitrary fashion. For example, something is 'grue' iff either (a) it is both green and observed before time t , or (b) it is both blue and observed after time t .

class of properties whose presence we must be able to distinguish from the presence of the original determinate property when we seek to non-luckily isolate the latter in thought. 4 follows once we admit these nearby-but-discriminable properties—those that share a subset of their conditional causal powers with the ascribed property. For when we try to isolate a property F within a diverse range of potential causal interactions, there will be rationally-relevant potential causal interactions involving the nearby-but-discriminable-from-F properties.

Objection and Reply The way of I've established 4 opens up a potential objection to my second argument: while ascription of a special causal relation requires isolating a type of change characteristic of that relation, an opponent might insist that this does not require a subject to isolate causally-fine-grained properties (in particular, a subject need not distinguish a property from others that would make the same contribution to the rationally-relevant causal interactions). But if ascribing a special causal relation merely requires ascribing causally-coarse-grained properties, we arguably lose the threat to our non-lucky isolation of these properties in thought. For suddenly any property we isolate using a set of conditional causal powers will count as an instance of the property we seek to isolate in thought.

I think we can resist this objection. When we ascribe causal relations—special or otherwise—we standardly treat them as immune to certain sorts of fragility.²³ Suppose I throw a ball that breaks a window. Had I thrown the ball in a radically different way, or had I thrown a much lighter ball, or had I thrown it at a time when the window was not there, my throw would not have broken the window. Thus certain alternations will turn an instance of a causal relation into a scenario where that causal relation does not obtain. However not every alternation has such a destructive effect on a causal interaction. In this sense causation isn't fragile. For instance, if I threw

²³ Cf. Lewis (2000).

the ball a moment earlier, or if the ball were an ounce heavier, or if the window were ten inches closer, my throwing of the ball would still break the window.

When we ascribe a special causal relation we must recognise this non-fragility of causation. After all, one would surely look askance at a subject who claims to be able to ascribe a special causal relation, but who refuses to ascribe the relation except when conditions exactly match those under which the relation was first ascribed. Yet if our isolation of the change characteristic of a special causal relation depends upon our ascription of causally-coarse-grained properties, we lose our cognitive grip on this non-fragility of causation. When I throw the ball a different way, my breaking of the window depends upon the presence of different properties. But for my throw to still involve the same causal relation, it must involve the same type of change (since we've typed special causal relations in terms of the kinds of change they involve). As a result, we can recognise the non-fragility of special causal relations only if our grasp of these changes rests on our ascription of properties as causally-fine-grained. For if we ascribe these properties in a causally-coarse-grained fashion, we'll lose the distinction between the original properties involved in a causal interaction, and the properties involved after causation-sustaining alterations.

In summary: to recognise special causal relations as non-fragile, we must ascribe properties as causally-fine-grained; but (as we saw above) if ascription of special causal relations rests on our ascription of properties as causally-fine-grained, we cannot simply ascribe these properties as dispositional—we must ascribe them as categorical.

4 Conclusion

Let's briefly review the ground I've covered. The central burden of this chapter has been to defend the following argument for PREDICATIVE ATOMISM:

1. CAUSALISM: Our capacity to ascribe unobservable properties exploits our capacity to ascribe the causal relations that hold in virtue of these properties (i.e. causal relations that constitute the causal roles of these properties).
2. CAUSAL ATOMISM: Our capacity to ascribe causal relations that hold in virtue of unobservable properties type-depends upon our capacity to ascribe special causal relations.

1 and 2 together entail 3:

3. Our capacity to ascribe unobservable properties type-depends upon our capacity to ascribe special causal relations. [From 1 and 2]
4. CONJECTURE: Our capacity to ascribe special causal relations type-depends upon our capacity to perceptually demonstrate objects (i.e. to identify objects solely on the basis of perceptual encounters with them).

As we saw in §3.1, CONJECTURE is John Campbell's answer to a puzzle Sydney Shoemaker raises for our conception of ordinary objects. In order to bolster Campbell's answer, I showed how a potential alternative answer to Shoemaker's puzzle—an answer that appeals to Lewis's notion of a 'natural property'—falls short. Given 3 and 4, 5 follows:

5. Our capacity to ascribe unobservable properties type-depends upon our capacity to perceptually demonstrate objects. [From 3 and 4]

In §3.2 I used 5 as part of an argument for 6:

6. BRIDGE PREMISS: Our capacity to ascribe special causal relations also type-depends upon our capacity to perceptually ascribe observable properties as categorical.

And 3 and 6 together entail

7. PREDICATIVE ATOMISM: our capacity to ascribe unobservable properties type-depends upon our capacity to perceptually ascribe observable properties as categorical. [From 3 and 6]

My main contribution to this argument for PREDICATIVE ATOMISM was my defence of the BRIDGE PREMISS in §3.2. I argued for the BRIDGE PREMISS by defending three nested component claims: (1)

our capacity to ascribe special causal relations type-depends upon our capacity to ascribe observable properties; (2) this ascription must be perceptual; and (3) we must ascribe at least some of the properties as categorical.

5 Looking Ahead: Perception and Thought

This final section looks ahead to my final two chapters. Each will develop a consequence of my argument for PREDICATIVE ATOMISM. I briefly sketch these consequences below.

My argument for PREDICATIVE ATOMISM places substantial constraints on a right account of perception. PREDICATIVE ATOMISM carves out a rich explanatory role for perceptual ascription: perceptual ascription of observable properties as categorical grounds our capacity to ascribe unobservable properties. As a result, perceptual ascription requires a perceptual encounter with the ascribed property. In these encounters perceptual experience must present the property ascribed. The perceptual ascription invoked by PREDICATIVE ATOMISM requires a special sort of perceptual presentation. But some accounts of perceptual experience cannot admit this sort of perceptual presentation—at least not without running afoul of the connection I’ve established between perceptual ascription and our grasp of causation. I shall argue in Chapter 3 that PHENOMENAL INTERNALISM cannot be correct if we admit the relevant perceptual presentation. PHENOMENAL INTERNALISM holds that subjects may differ with respect to the objects and properties in their environment, and in their relationships to these entities, yet nevertheless what it is like for them when perceiving may be the same. I turn the tension between PREDICATIVE ATOMISM and PHENOMENAL INTERNALISM into an argument for the claim that what it is like to undergo a perceptual experience depends upon facts about which observable properties an experience presents.

My argument for PREDICATIVE ATOMISM also undermines a powerful challenge, now largely forgotten, to the claim that ascribing a property—as one does when perceptually ascribing it—

counts as a way of thinking about it. Dummett (1973/1981, Ch. 7), to whom we owe the original challenge, assumes that we can isolate the explanatory role of property ascription just by considering what is involved in grasping the truth conditions of sentences such as ‘Duncan is tall’. He uses this assumption to motivate a sharp distinction between requirements on thinking about objects and requirements on the ascription of properties.

Dummett says that to think about an object we must somehow single it out (e.g. perception enables you to distinguish a ball from other objects). What does it take to single out an object? Perhaps you must be able to track the object through change, or to re-identify it after gaps in observation, or maybe you need to know persistence conditions characteristic of objects of its kind. Property ascription is different, insists Dummett; ascription of a property requires only that you be able to recognise instances of it. And this recognitional capacity seems to require only sensitivity to the property’s instantiation—not any richer operation of singling out. Plausibly, when you recognise something as satisfying ‘cat’ or ‘dog’, for example, you neither keep track of those kinds (properties) nor exploit specialized knowledge of the conditions for being a cat or being a dog. Thus on Dummett’s view, ascribing a property is not like thinking about an object; indeed it’s not a way of *thinking about something* at all.

The paradigm cases of ways of thinking about entities are the ways we think about objects. Dummett has put his finger on an apparently deep contrast between thinking about objects and ascribing properties. Chapter 4 resurrects a more general version of Dummett’s challenge: I show that Dummett’s contrast threatens the widespread assumption that property ascription is a genuine way of thinking about a property. Later I show that my argument for PREDICATIVE ATOMISM generates the beginnings of a response to the threat posed by Dummett’s challenge. My argument provides an explanatory role for perceptual ascription that motivates richer requirements on the perceptual ascription of properties. These requirements go beyond mere sensitivity to the

instantiation of a property, and license the claim that perceptual ascription is a way of thinking about a property after all.

Chapter 3

Whence Come Properties Before The Mind?

In this chapter I show how my argument for PREDICATIVE ATOMISM constrains a right account of perceptual experience.

To get PREDICATIVE ATOMISM back onto the table—and to bring the project of this chapter into view—I shall briefly remind the reader of the framework developed in the last chapter. Recall that a subject can ascribe a property only if she knows what it would be for something to have the property; this knowledge in turn rests on a kind of classificatory knowledge. Classificatory knowledge is manifested by a subject’s sensitivity to conditions under which a property counts as instantiated.

The framework of Chapter 2 types ways of ascribing properties by the kind of classificatory knowledge they involve. To ascribe a property ‘as dispositional’ a subject must possess classificatory knowledge limited to knowledge of causal powers that the property confers upon objects that instantiate it. By contrast, to ascribe a property ‘as categorical’ she must possess classificatory knowledge that satisfies two conditions: (1) it cannot be limited to knowledge of causal powers that the ascribed property confers; and (2) it must justify treating the ascribed property as a ground of an object’s causal dispositions (for instance by enabling a subject to recognise that round objects roll *because* they are round). Understood within this framework, PREDICATIVE ATOMISM is the thesis that our capacity to ascribe unobservable properties depends upon our capacity to perceptually ascribe observable properties as categorical.

What must perceptual experience be like for it to enable a subject to perceptually ascribe an observable property as categorical? Within my framework this becomes a question about the epistemic reach of perceptual experience: what must perceptual experience be like in order for it to

help provide the classificatory knowledge required for perceptual ascription of observable properties as categorical? Following Russell (1910), call this the question of what perceptual experience must be like in order to place a property ‘before the mind’.¹

I shall argue that a large class of extant accounts of perceptual experience (many of them widely held) cannot permit perceptual experience to provide the classificatory knowledge required by our capacity to perceptually ascribe observable properties as categorical. These accounts presuppose PHENOMENAL INTERNALISM—the view that subjects may differ with respect to the objects and properties in their environment, and in their relationships to these entities, yet nevertheless what it is like for them when undergoing a perceptual experience may be the same.² (This is often expressed as the view on which what it is like to undergo a perceptual experience supervenes upon a subject’s intrinsic properties.)

The chapter has five parts. §1 introduces some important distinctions from the literature on perceptual experience. §2 uses PREDICATIVE ATOMISM to motivate a substantive constraint on a right account of perceptual experience. §3 reconstructs a powerful argument from David Chalmers and Brad Thompson. The argument begins with PHENOMENAL INTERNALISM, but its conclusion conflicts with the constraint motivated in §2. In §4 I canvas ways to avoid the conflict without sacrificing PHENOMENAL INTERNALISM; none of these escape routes proves very attractive. §5 uses this result to motivate a version of PHENOMENAL EXTERNALISM on which perceptual experience places us in direct contact with observable properties.

¹ Johnston (2004) and Pautz (2010) exploit different versions of my question to argue against competing accounts of visual experience: put roughly—and abstracting from the differences between them—they ask what hallucinatory visual experience must be like to permit subjects to think about hallucinated properties. I’m not inclined to grant their presupposition that hallucinatory visual experience permits thought about hallucinated properties; I’m even less convinced by their attempts to exploit this presupposition in order to establish substantive conclusions about visual experience. An analogue of my question also arises for perceptual demonstrative thought about objects (cf. Campbell 2002; Smithies 2011; Dickie 2015): what must visual experience be like in order to permit perceptual demonstrative thought about perceived objects?

² Recent defenders of versions of PHENOMENAL INTERNALISM include (to name just a few) Shoemaker (1994); Thau (2002); Chalmers (2004, 2006a); Pautz (2006).

1 Conceptual Topography

Suppose you're half way up a climbing wall. What is it like for you, perceptually speaking? Probably something like this: you feel your hands beginning to sweat, despite the liberal application of powdered chalk; your attention is fixed on the textured handhold lying just beyond your reach; and your breathing sounds loud in your ears. This specification of what it is like for you is also a specification of the 'phenomenal character' of your perceptual experience:

Phenomenal Character The phenomenal character of a perceptual experience is what it is like for a subject to undergo the experience.

Philosophers standardly type perceptual experiences according to their phenomenal character.³ Any two experiences with a given phenomenal character will be typed together. So if I'm climbing the same wall, but on a different day, and what it is like for me at the half-way point happens to be the same as before, my perceptual experience will count as a token of the same type as my earlier experience.

Now 'what it is like to undergo a perceptual experience' may cover diverse aspects of perceptual experience: some presentational (roughly: which bits of the world the experience seems to reveal); others affective (roughly: how an experience *feels*); and still other⁴ aspects that fall under neither of these categories. I take no stand on whether there are non-presentational aspects of perceptual experience, or on what those aspects might be. (Some⁵ go so far as to deny that perceptual experience ever has non-presentational phenomenal character.) My interest instead lies with the presentational phenomenal character of perceptual experience. The remainder of this

³ Of course, there are other ways to type experiences: we can type them by when they occur, in whom they occur, which judgments they dispose subjects to make, and so on. Yet typing them by their phenomenal character seems to get at what fundamentally distinguishes one conscious experience from another. See Soteriou (2005, p. 194) for discussion of this point.

⁴ Peacocke (1983), for example, argues that a right account of perceptual experience must attribute non-presentational non-affective 'sensational properties' (sometimes called 'qualia') to perceptual experience.

⁵ Dretske (1995); Tye (1995; 2000).

section introduces distinctions within and around presentational phenomenal character: §1.1 says a bit more to isolate this sort of phenomenal character; §1.2 sketches two competing accounts of what it is for a perceptual experience to present bits of the world; §1.3 introduces the notion of ‘phenomenal content’.

1.1 Perceiving the World

Even untutored reflection suggests that (at least) part of what it is like to undergo a perceptual experience is to have certain bits of the world—objects, colours, shapes, and the like—perceptually presented. As P.F. Strawson (1979) observes:

Suppose a non-philosophical observer gazing idly through a window. To him we address the request, “Give us a description of your current visual experience,” or “How is it with you, visually, at the moment?” Uncautioned as to exactly what we want, he might reply in some such terms as these: “I see the red light of the setting sun filtering through the black and thickly clustered branches of the elms; I see the dappled deer grazing in groups on the vivid green grass...” and so on.⁶

As I’ll use the term, ‘presentational phenomenal character’ captures both what an experience seems to present—the deer, the elms, their respective colours—and how these elements of the world show up from a subject’s perceptual perspective. When I see Strawson’s dappled deer, for instance, I see them from a particular angle, under certain lighting conditions, with my visual systems working a certain way, and so on. Were I to see the same deer from a different perspective, the presentational phenomenal character of my experience might be quite different: I might see different portions of the deer, or their colour might look a bit different, or my vision might be so blurry that I’m no longer in a position to recognise them as deer at all.

⁶ Strawson (1979, p. 93)

1.2 Varieties of Presentation

Fierce disagreement arises when philosophers try to say something more substantive about presentational phenomenal character. In particular, I've used the placeholder notion 'presentation' to isolate presentational phenomenal character. Yet what is it for an experience to render something *present* to the mind?

Contemporary accounts of perceptual presentation mostly fall into two broad camps: representationalism and direct realism.⁷ Put briefly: representationalists explain perceptual presentation as a variety of *representation*, whereas direct realists explain it as a *non-representational relation*. I'll provide more detail about these two views in just a moment. But let me first explain why I mention them at all: I need them in order to show that the notion of phenomenal content—a notion not yet introduced—remains neutral between these competing views.

I'll begin with representationalism. Its central assumption is that perceptual experience possesses a variety of representational content. As we saw in Chapter 1, a standard way for philosophers to motivate the attribution of representational content to perceptual experience is by considering cases that seem to involve perceptual error: a white object under red illumination looks red; a straight stick partially immersed in water looks bent; a round vase viewed through distorting lenses looks oval. One way to describe these cases is as involving a 'mismatch' between what experience presents (the way the world 'looks'), and what the world is actually like. Many introduce representational content to capture the (supposed) distinction between 'world-matching' and 'world-mismatching' perceptual experience. A perceptual experience 'matches' how things stand in the perceived scene only if it accurately represents the scene. Thus an experience's 'representational

⁷ Recent representationalists include Harman (1990); Shoemaker (1994); Dretske (1995); Tye (1995, 2000); Lycan (1996); Siewert (1998); Byrne (2001); Thau (2002); Levine (2003); Chalmers (2004, 2006a); Thompson (2010). Direct realists—an increasingly influential minority amongst philosophers of perception—include Campbell (2002); Martin (2002); Fish (2009); Hellie (2010); Brewer (2011); Logue (2012); Soteriou (2013). Of course, there are also those who try to keep a foot in both camps, such as Schellenberg (2010).

'content' is the way the world must be if it is to 'match' the experience (or, in more jargon-y terms, it is the way the world must be if the experience is *veridical*). Given this background we can nail down our definition of representationalism as follows:

REPRESENTATIONALISM The presentational phenomenal character of a perceptual experience is exhausted by the experience's representational content.⁸

It helps to see the view in action. Suppose I undergo a visual experience as of a red cup on my table. My experience will be veridical (that is, things will be as they look) only if there is a red cup on my table. So the representational content of my experience will be something like this: *there is a red cup on the table*. Given this representational content, representationalism says that the presentational phenomenal character of my experience amounts to this: *there is a red cup on the table*.

Representationalists will say that a perceptual experience presents an object or property just if the entity belongs to the experience's representational content. Thus all perceptual presentation is a matter of representation (and the *ways* of being presented with an object or property—those that reflect a subject's perceptual perspective—are just the ways an experience represents these entities).

Direct realism is the main contemporary alternative to representationalism.⁹ Unlike representationalists, direct realists initially restrict themselves to the phenomenal character of veridical experience (though they do not explain veridicality in terms of a 'match' between an experience and the world). As a result, many internal disputes between direct realists concern how to extend the account to cover illusory and hallucinatory experience.

⁸ Representationalists who wish to *reduce* phenomenal character to representational content will usually make the stronger claim that an experience's representational content exhausts every dimension of phenomenal character, not just presentational phenomenal character (Cf. Tye 2000).

⁹ The view goes by many names (abstracting from differences irrelevant to the present discussion): 'direct realism' (Hellie 2010), 'naïve realism' (Martin 2002), 'the object view' (Brewer 2011), and 'the relational view' (Campbell 2002; Soteriou 2013).

DIRECT REALISM The presentational phenomenal character of a veridical perceptual experience consists in the external world objects and properties the experience presents, together with the perceiving subject's perspective on those objects and properties.

The perceptual relation a subject bears to these entities isn't representational. We do not perceive *that* an object has certain properties. We instead perceive the object *and* its properties. Presentation becomes relational rather than representational.

An example will help bring out the contrast between direct realism and representationalism. Suppose you see a black cat sitting on your porch. Direct realism says that the phenomenal character of your visual experience of the cat consists in (a) the cat itself, (b) its visible properties (its black colour, its sleek shape, etc.), and (c) the various elements of your perceptual perspective on the cat (your location, where you've directed your perceptual attention, the lighting conditions, how your eyes are functioning, etc.). If any element from (a)–(c) were to change, the phenomenal character of your experience would also change. For instance, you might see a different black cat, or the same cat with different visible properties. Or you might change your perspective on the cat by walking around to view it from another angle, or shift your attention from its ears to its tail.

More generally, direct realism entails that two veridical perceptual experiences have the same presentational phenomenal character only if both (1) the subjects perceive the exact same objects and properties, and (2) the relevant perspectival factors are identical. Again, the relation a subject bears to these entities isn't representational—we do not perceive *that* an object has certain properties; we instead perceive the object *and* its properties.

Stepping back a bit, what remains invariant between representationalism and direct realism is a concern to capture the fact that perceptual experience possesses what I've called 'presentational phenomenal character': perceptual experience seems to reveal objects and their observable properties (and does so from a particular perspective). The two views simply provide different accounts of how perceptual presentation works.

1.3 Phenomenal Content and Perceptual Presentation

Now that we're armed with a relatively neutral notion of presentational phenomenal character, I shall use it to define yet another notion: 'phenomenal content'. Put intuitively, the phenomenal content of a perceptual experience is what the experience presents in virtue of having the presentational phenomenal character that it does. More carefully:

Phenomenal Content The phenomenal content of a perceptual experience consists of what, necessarily, two experiences both present if they have the same presentational phenomenal character.¹⁰

This definition remains neutral between representationalism and direct realism. On representationalism, what an experience presents is captured by a kind of representational content (namely that representational content which two experiences must share if they possess the same presentational phenomenal character). Direct realism, by contrast, does not use representational content; perceptual experience is instead constituted by the perceptual presentation of objects and their properties from a particular perceptual perspective. The phenomenal content of an experience therefore consists of the objects, properties, and perspectival factors (if any) that must be presented (or simply present, in the case of perspectival factors) whenever a subject undergoes an experience with the same presentational phenomenal character as the original.

An example will consolidate this otherwise abstruse notion of phenomenal content. I introduced phenomenal character by providing a rough and ready description of what it is like for you when half way up a climbing wall: your hands are beginning to sweat; your attention is fixed on the handhold just beyond your reach; and so on. If this rough description is an accurate guide to the phenomenal character of your experience, it will also be a guide to what the experience presents in virtue of its presentational phenomenal character. A representationalist will say: the phenomenal

¹⁰ This definition generalises the notion of phenomenal content given in Chalmers (2006a, p. 50–1). He attributes the primary argument for phenomenal content to Siewert (1998).

content of your experience—the representational content it possesses in virtue of its presentational phenomenal character—is that your hands are beginning to sweat, and that your attention is fixed on a handhold just beyond your head, and so on. A direct realist will say: the phenomenal content of your experience—the objects and properties the experience presents (from a given perspective) in virtue of its presentational phenomenal character—consists of your hands, their sweatiness, and the nearby handhold (where your perspective on the last involves attentional focus).

One obvious difference between the representationalist and direct realist stories concerns the relationship between what is presented and what exists in the perceived scene. A representationalist can allow my experience to represent that my hands are sweating, when in fact my hands aren't sweating. Yet a direct realist cannot allow the phenomenal content of my experience to include a property (like sweating) unless I actually perceive something that has it.

The introduction of phenomenal content may seem to be in tension with the argument of Chapter 1. For in that chapter I argued against a standard explanation of perception-based rational belief formation—what I called the ‘tempting explanation’. A central feature of the tempting explanation was its appeal to a type of perceptual content—a type of content capable of being accurate or inaccurate—motivated by the possibility of cases in which objects look (or sound, or feel...) other than they are (such as when a white table under red illumination looks red). The tempting explanation used this sort of perceptual content to explain which beliefs a subject would be rational to form in response to an experience. It may have appeared that my argument against the tempting explanation threatens any attribution of content to perceptual experience. However this assumption does not withstand scrutiny. The tempting explanation runs aground because of its *use* of such content to explain rational belief formation, not the attribution of content as such.

Furthermore, phenomenal content (as I've defined it) need not occupy the explanatory role assigned to perceptual content by the tempting explanation. On direct realism phenomenal content is neither accurate nor inaccurate—unlike the content required by the tempting explanation. For as

I've already shown, to specify phenomenal content a direct realist need only identify the objects and properties presented by an experience, as well as the perspective from which they are presented. (However direct realists might still permit an explanatorily derivative type of accuracy-conditioned content: there are perceptual perspectives on an object such that beliefs a subject would form by uptake from perception come out true; one might treat an experience of the object as 'accurate' iff it involves a perceptual perspective of this sort.) According to representationalism, by contrast, phenomenal content is a kind of representational content. It is therefore capable of being accurate or inaccurate, as the tempting explanation requires. Yet proponents of representationalism can dodge my argument against the tempting explanation if they sever the tempting explanation's link between perceptual content and rational belief formation. So both direct realists and representationalists can admit phenomenal content without fear of running afoul of my argument against the tempting explanation.

2 How PREDICATIVE ATOMISM Constrains Perception

This section uses PREDICATIVE ATOMISM to motivate a particular constraint on the phenomenal content of perceptual experience: namely, this content must include those observable properties we can perceptually ascribe as categorical. §2.1 traces the path from PREDICATIVE ATOMISM to the constraint; §2.2 defends one of the central steps of the argument of §2.1.

2.1 The Path from PREDICATIVE ATOMISM

The definition of perceptual ascription of properties as categorical entails 1:

1. Ascription of a property as categorical requires classificatory knowledge that goes beyond knowledge of the causal role occupied by the ascribed property.

If we take on board the anti-sceptical assumption that we do manage to ascribe unobservable physical properties and causal relations, PREDICATIVE ATOMISM entails that

2. We can perceptually ascribe observable properties as categorical.

But in §2.2 I shall argue that

3. Only the phenomenal content of perceptual experience could provide the classificatory knowledge required by this type of perceptual ascription.

So PREDICATIVE ATOMISM constrains the phenomenal content of perceptual experience:

4. The phenomenal content of perceptual experience must provide the classificatory knowledge required to perceptually ascribe observable properties as categorical. [From 1–3]

Those who accept 4 must tell a story about how the phenomenal content of perceptual experience provides the requisite classificatory knowledge. How might this story go? Given my definition of ‘phenomenal content’, any special ability on the part of this content to provide classificatory knowledge will surely rest on what constitutes the content (namely what an experience with that content presents in virtue of its presentational phenomenal character). However if the phenomenal content does not include observable properties—in other words if perceptual experience does not present the properties themselves—it becomes quite hard to understand how phenomenal content could provide the classificatory knowledge required by 4. If this is right, 4 leads to 5:

5. Perceptual experience must present the observable properties we can perceptually ascribe as categorical (that is, these properties must figure in the phenomenal content of perceptual experience). [From 4]

To see why 5 likely follows from 4, suppose that the presentational phenomenal character of my experience of a red cup requires only the perceptual presentation of a higher order property of red—say that of being the property whose instantiation normally causes me to undergo visual experiences with a ‘reddish’ phenomenal character—rather than to red itself.¹¹ Perceptual presentation of such a property plausibly cannot provide classificatory knowledge of red that both (a) goes beyond

¹¹ In §3 we’ll see David Chalmers and Brad Thompson endorse something quite like this claim about phenomenal content. They argue that if PHENOMENAL INTERNALISM is true, the phenomenal content of perceptual experience cannot include observable properties (a result clearly inconsistent with 5).

knowledge of the causal powers red confers upon red objects and (b) justifies treating red as a ground of an object's causal dispositions. Put intuitively: it only permits us to know red as *whichever property possesses that higher order property*; it doesn't permit us to know red as such.

I have not given a story about how phenomenal content provides the classificatory knowledge that underwrites our capacity to perceptually ascribe observable properties as categorical. And without such a story in place, I cannot rule out the possibility of an alternative on which 4 does not entail 5. The transition from 4 to 5 therefore represents a point at which an opponent may resist my attempt to derive some claims about perception from a claim about thought. However I have said enough to make the transition from 4 to 5 difficult to dismiss.

Let's take stock for a moment. I've given an argument that traces the path from PREDICATIVE ATOMISM to 5 (the claim that perceptual experience must present the observable properties we can perceptually ascribe as categorical). 1 follows from the framework of Chapter 2. PREDICATIVE ATOMISM gives us 2. And once we have an argument for 3, 1–3 give us 4. I've already said something in defence of the final move from 4 to 5. So what remains to be done is to defend 3.

2.2 Basic Cognitive Capacities and Immediate Justification

This subsection defends 3—the claim that only the phenomenal content of perceptual experience could provide the classificatory knowledge required by our capacity to perceptually ascribe observable properties as categorical. I first argue that this capacity is *cognitively basic*; I then show how that claim leads to 3.

2.2.1 Our capacity to perceptually ascribe observable properties as categorical is cognitively basic

A cognitive capacity is 'basic' in my sense iff it does not depend upon other cognitive capacities belonging to the same kind. On this definition our capacity for perceptual demonstrative thought likely qualifies as cognitively basic (specifically, it likely qualifies as a basic way of identifying objects).

Suppose a subject thinks *that is red* of a perceived object; must she also be able to identify objects some other way? Surely not: if anything constitutes our most primitive way of identifying objects in the world, it is perceptual demonstration. ATOMISM provides a deeper anchor for this claim. Introduced at the beginning of Chapter 2, ATOMISM is the view that our capacity for descriptive identification of objects depends upon our capacity for perceptual demonstrative thought about ordinary objects. Since descriptive identification and perceptual demonstration likely exhaust the candidates for our most basic mode of identification of objects, ATOMISM helps explain why our capacity for perceptual demonstrative thought must be a basic cognitive capacity.¹²

When I originally introduced PREDICATIVE ATOMISM, it was as an analogue of ATOMISM. Given this parallel, we should expect PREDICATIVE ATOMISM to strongly support the (already intuitively plausible) claim that our capacity for perceptual ascription of observable properties as categorical is a basic cognitive capacity.

Whether a particular type of property ascription counts as cognitively basic will depend upon the classificatory knowledge that underwrites it. If that knowledge does not rest on some other classificatory knowledge, the property ascription will be cognitively basic. PREDICATIVE ATOMISM tells us that the classificatory knowledge that underwrites ascription of unobservable properties ultimately depends upon the classificatory knowledge that underwrites perceptual ascription of observable properties as categorical; and my argument for PREDICATIVE ATOMISM tells us that this classificatory knowledge also grounds the classificatory knowledge involved in our ascription of causal relations. With these two sorts of classificatory knowledge depending upon the classificatory

¹² This claim does not presuppose that descriptive identification and perceptual demonstration are the only ways we have to identify objects. For instance, I might exploit my grasp of a proper name—say ‘Psyche’—to think about someone I’ve never met, and about whom I know nothing (except her name). Having never met Psyche, I’m not in a position to identify her demonstratively; yet without knowing any further facts about her, it is unlikely that I’m able to identify her descriptively (cf. Kripke 1980). However even if proper-name-based thought involves a mode of identification that is neither demonstrative nor descriptive, it would be highly implausible to insist that our capacity for such thought instead underwrites perceptual demonstrative thought (cf. Evans 1982, pp. 373–5).

knowledge characteristic of perceptual ascription of properties as categorical, what could the latter classificatory knowledge possibly depend upon? I see no plausible alternatives. PREDICATIVE ATOMISM therefore strongly suggests that perceptual ascription of observable properties as categorical ranks among our basic cognitive capacities.

2.2.2 Classificatory Knowledge and Phenomenal Content

I've argued that our capacity to perceptually ascribe observable properties as categorical is cognitively basic. I shall now show how this claim leads to 3: the claim that only the phenomenal content of perceptual experience could provide the classificatory knowledge required by our capacity to perceptually ascribe observable properties as categorical

Property ascription always involves classificatory knowledge. Perception helps secure this knowledge when we *perceptually* ascribe properties. But on its own this tells us little about what role perception must play: perception might play a merely causal role in securing the requisite classificatory knowledge. For example, the epistemic role of perception might be exhausted by its status as a process that reliably causes subjects to form true beliefs about the world.¹³ Yet such an austere causal story would likely abandon a highly plausible claim about how perceptual ascription works: perception enables perceptual ascription of *this* property, as opposed to ascription of *that* property, in part because of which chunk of the world it presents. So we should expect a close tie between what perception presents and which properties it permits us to perceptually ascribe (and which ways we are permitted to ascribe them).

Now even if perception plays more than a merely causal role, some will insist that its contribution turns on which other beliefs a subject has about the world. To see how this might work, consider select cases where perception secures non-classificatory knowledge about the world.

¹³ A view of this sort would count as a version of ‘process reliabilism’. For a seminal discussion of process reliabilism, see Goldman (1979).

If I see gnocchi rise to the surface in a pot of boiling water, I know it is done; but I know this only because I'm already aware of what it takes for gnocchi to cook properly. In this case, perception justifies my belief—and so bumps it up to the level of knowledge—only in virtue of my prior justified beliefs about what it takes for gnocchi to cook. Call this a case of ‘mediate justification’ by perception (the case of Dorothy in Chapter 1 was a case of this sort).

Those who agree that perception plays more than a purely causal role in perceptual ascription, yet who still wish to minimize the epistemic burden carried by perception, might try to insist that perception secures classificatory knowledge only through mediate justification. For instance, suppose I know some general law-like connection between the instantiation of a certain property in the world, and my undergoing a perceptual experience with such-and-such presentational phenomenal character. Perhaps I know that when I feel my desk vibrating, it is because a subway train is passing through a tunnel that runs underneath the building. I could combine this knowledge with my knowledge of what it is like to undergo a desk-vibrating experience, and thereby come to possess classificatory knowledge of a property of my building—in this case the property of having a subway train pass close underneath.

This ‘mediate justification’ story becomes more difficult to sustain when we turn to types of perceptual ascription that count as cognitively basic. If the classificatory knowledge that underwrites such perceptual ascription were to involve mediate justification, this justification would depend upon a subject's justification for other beliefs. These beliefs would also involve a type of property ascription. And this ascription would in turn rest upon yet more classificatory knowledge. So it looks as though a capacity to engage in perceptual ascription will count as cognitively basic only if the underlying classificatory knowledge is not the product of mediate justification by perception. The required non-mediate justification—what epistemologists have come to call ‘immediate

'justification'—is justification that does not depend upon justification for any of a subject's other beliefs about the world.¹⁴

We've now come quite close to 3—the claim that only the phenomenal content of perceptual experience could provide the classificatory knowledge required by our capacity to perceptually ascribe observable properties as categorical. I argued in §2.1.1 that this capacity is cognitively basic. And I've now argued that the status of this capacity as cognitively basic constrains the epistemic role of perceptual experience. Perception must provide immediate justification for beliefs that involve perceptual ascription of a property as categorical (beliefs like *it's that shape* or *this size is larger than that size*).

In order to reach 3 we require a final claim: the claim that the immediate justification perceptual experience provides must wholly depend upon its phenomenal content (in other words, perceptual experiences with the same phenomenal content will provide the same immediate justification).¹⁵

A full defence of this claim would require a dissertation of its own. But thankfully the argument of this chapter does not require a worked-out defence of the claim. My goal is to use PREDICATIVE ATOMISM to exert pressure on PHENOMENAL INTERNALISM. It's therefore enough to show that commitment to the combination of PREDICATIVE ATOMISM and PHENOMENAL INTERNALISM generates real costs. I merely need to say enough to show that separating immediate justification and phenomenal content constitutes a serious cost.

I'll begin with a reason to avoid loosening the restriction of immediate justification to phenomenal content. Consider what would become possible were we to lift the restriction. Suddenly

¹⁴ Cf. Pryor (2000, p. 532–5) for a representative discussion of 'immediate justification'.

¹⁵ Silins (2011, p. 335) dubs a representationalist version of this claim the 'Content Constraint'. He argues against it, which pits him against many of those who accept that perception can provide immediate justification. However he also shows that the constraint underlies a powerful version of the old (yet influential) 'veil of perception' arguments against indirect realism about perception. (Indirect realism is the view that the immediate objects of perceptual awareness are mind-dependent entities, rather than mind-independent objects and their properties.)

two perceptual experiences with the *same* phenomenal content (i.e. what these experiences present in virtue of their presentational phenomenal character is the same) might nevertheless immediately justify *different* beliefs. Yet this result seems absurd. Imagine what it looks like if you're a representationalist. Two subjects—Anna and Vronsky—might undergo experiences with the same phenomenal content; perhaps the representational content *there is a blue ball*. In this case part of what it is like for each of them (in particular, their presentational phenomenal character) is that there is a blue ball before them. Yet nevertheless there will be beliefs such that if Anna formed one of them in response to *her* experience, it would be immediately justified; but if Vronsky formed the same belief (namely a belief with the same content as Anna's) in response to *his* experience, the belief would not be immediately justified. The situation remains strange even when we switch from representationalism to direct realism. A direct realist will say that Anna and Vronsky share a phenomenal content in virtue of sharing a perceptual perspective upon a particular object and its colour—in this case the ball and its blueness. Yet were direct realists to lift the restriction of immediate justification to phenomenal content, they would have to allow that the perceptual experiences of Anna and Vronsky provide immediate justification for different beliefs (despite the sameness of their perceptual perspectives upon an object and its colour). And again, this result seems quite bizarre.

We also have positive reason to accept the restriction of immediate justification to phenomenal content: the restriction helps explain intuitive boundaries on immediate justification. Immediate justification tracks what a subject would have grounds to believe in response to perceptual experience if she were to bracket her other beliefs about the world. Which beliefs satisfy this condition? A natural way of finding out is to ask Strawson's question: 'How is it with you, visually, at the moment?' If we're inclined to restrict the phenomenal content of visual experience to low-level features like shapes and colours and the like, we might ask a more emphatic version ('How is it with you, *visually*, at the moment?'). In any case, it seems that we often find ourselves collapsing

questions about the reach of immediate justification into questions about presentational phenomenal character. This pattern becomes easy to explain if immediate justification wholly rests upon phenomenal content.

We've reached the end of another section. I've used PREDICATIVE ATOMISM to argue for the following constraint on perception (5 from the §2.1): perceptual experience must present the observable properties we can perceptually ascribe as categorical (that is, these properties must figure in the phenomenal content of perceptual experience—regardless of whether we are representationalists or direct realists). In the next section I argue that this constraint is in tension with PHENOMENAL INTERNALISM. Given my argument for PREDICATIVE ATOMISM in Chapter 2, the tension undermines PHENOMENAL INTERNALISM.

3 An Internalist Wedge

My argument against PHENOMENAL INTERNALISM turns on the claim that no defensible version of PHENOMENAL INTERNALISM can admit ordinary observable properties into the phenomenal content of perceptual experience. The task of this section is to set out a powerful argument in support of this claim—an argument developed by David Chalmers and Brad Thompson¹⁶—and briefly sketch some ways for phenomenal internalists to resist its conclusion. (I postpone to §4 a more thorough investigation of whether these ways of resisting the argument can save PHENOMENAL INTERNALISM from PREDICATIVE ATOMISM.)

A range of accounts of perceptual experience share a commitment to PHENOMENAL INTERNALISM:¹⁷

¹⁶ They develop versions of the argument piecemeal for different properties, with only the suggestion that the line of thought generalises. For the application to spatial properties, see Thompson (2010) and Chalmers (forthcoming); for colours, see Chalmers (2006a); and for temporal properties, see Chalmers (2012, p. 334).

¹⁷ Those committed to PHENOMENAL INTERNALISM are a varied bunch. For example, 'Fregean Representationalism' (cf. Chalmers 2004; Thompson 2010) says that the phenomenal content of perceptual experience consists of 'Fregean' contents like *there is an object with the property that normally causes experiences of this sort*. These contents are labelled 'Fregean'

PHENOMENAL INTERNALISM What it is like for a subject to undergo a perceptual experience supervenes upon the subject's intrinsic properties (and so does not depend upon which objects or properties are in their environment, or upon their relationships to these entities).

The Chalmers and Thompson argument begins with PHENOMENAL INTERNALISM, and ends up banishing observable properties from the phenomenal content of perceptual experience. Their argument runs as follows.

Someone could live in a world that systematically differs from the actual world (e.g. a 'stretched' world in which all distances are double those in the actual world)—a 'displaced scenario'—and nevertheless (given PHENOMENAL INTERNALISM) undergo experiences type-identical to our own (recall that for present purposes we are typing experiences by their phenomenal characters, in other words by what it is like for a subject to undergo them).¹⁸

Would these perceptual experiences be systematically in error? For such error to occur, there would have to be a mismatch between how the experiences present things as being and how things stand in the displaced scenario. For example, all objects would have to be three times as large as perceptual experience presents them being. Call this sort of systemic error 'total illusion'.

The possibility of total illusion seems to generate a nasty form of scepticism. If subjects in displaced scenarios whose experiences are type-identical to our own would be victims of total

because their accuracy conditions are not property-dependent. In a displaced scenario an experience with the same phenomenal content could turn out to be veridical, since 'the property that normally causes experiences of this sort' is a description satisfied by different properties in different worlds. However not every phenomenal internalist accepts Fregean Representationalism. For instance, someone could defend a version of 'sense datum theory' (cf. Jackson 1977)—a version on which perceptual experience involves direct awareness of mind-dependent entities—and still accept PHENOMENAL INTERNALISM. Apart from these two, many other views are compatible with—or entail—PHENOMENAL INTERNALISM.

¹⁸ Some will doubtless question the legitimacy of such inter-personal comparisons of phenomenal character. Yet someone could run a modified version of the Chalmers/Thompson argument using only diachronic intra-personal comparisons. A subject might live in a world that transforms into a displaced scenario too slowly and systematically for anyone to notice.

illusion, then (for all we know) we might actually be victims of total illusion. We take our perceptual experience to reveal the standard against which we compare the experiences of our experiential duplicates in displaced scenarios, but nothing ensures that we haven't things backwards. Their perceptual experiences could instead determine the standard against which we must hold our perceptual experiences (with us as the victims of total illusion).

Yet this kind of scepticism seems spurious. This anti-sceptical reaction mirrors an aspect of the puzzle raised by Shoemaker (1988) that CONJECTURE answered in Chapter 2. Shoemaker's puzzle arises when we ask what resources are required to capture our ordinary conception of physical objects; it acquires bite because natural answers to the question cannot distinguish ordinary objects from their 'shadow' counterparts (shadow counterparts are constructed by temporally displacing the set of space-time points that constitute the temporal trajectory of the original object). Just as the possibility of 'shadow' objects threatens to render our conception of physical object arbitrary, the possibility that we (or anyone else) could be victims of total illusion suggests that arbitrary prejudice supports our assumption that perception provides epistemic access to the external world. Shoemaker's puzzle concerns *the world* as we find it, while the present puzzle concerns *how* we find the world.

Quite apart from the spurious character of the scepticism induced by the possibility of total illusion, we also do not naturally regard people with experiences type-identical to our own, but who live in a displaced scenario, as victims of total illusion. But this natural judgment will be correct only if total illusion is impossible. Yet unless we rule out displaced scenarios on empirical or metaphysical grounds, only our account of perceptual experience can rule out the possibility of total illusion.

In order to rule out total illusion, Chalmers and Thompson restrict the phenomenal content of perceptual experience to features that remain invariant through displacement (for instance, the property of being the normal cause of an experience with such-and-such phenomenal character). On their 'Fregean Representationalism', perceptual experience contingently presents ordinary observable

properties as the bearers of the relevant invariant features. For example, in the actual world a perceptual experience of mine with a ‘reddish’ phenomenal character will present red as the actual normal cause of experiences with that phenomenal character. But in a displaced world experiences type-identical to my own will present green rather than red (assuming that the normal cause of ‘reddish’ perceptual experiences in that world is green).

We can represent the Chalmers/Thompson argument as follows (I stick with size properties for convenience, but the argument applies *mutatis mutandis* for other observable properties):

1. Displaced scenarios involving size properties—for example scenarios in which everything is systemically larger or smaller than in the actual world—are possible.
2. These scenarios may include subjects whose perceptual experiences are type-identical to our own. [From PHENOMENAL INTERNALISM]

So

3. If these subjects are victims of total illusion with respect to size properties, we face an invidious form of scepticism. [From 1 and 2]

But

4. This form of scepticism is spurious. Furthermore, subjects in these displaced scenarios whose perceptual experiences are type-identical to our own do not qualify as victims of total illusion.

So

5. Total illusions are impossible. [From 3 and 4]

Given the possibility of displaced scenarios involving size properties,

6. Our account of perceptual experience must explain why total illusions are impossible.
[From 1 and 5]

Yet 6 seems to require 7:

7. To explain why these illusions are impossible one must deny that size properties figure in the phenomenal content of perceptual experience.

And 6 and 7 entail 8:

8. The phenomenal content of perceptual experience does not include size properties.

[From 6 and 7]

Someone who accepts 8 can still posit a type of perceptual content that includes size properties (indeed Fregean Representationalists posit such additional content). But this content cannot be accorded the status of phenomenal content. In principle, two subjects might undergo experiences that differ with respect to this new level of content, but which nevertheless possess the same presentational phenomenal character.

I argued in §2 that PREDICATIVE ATOMISM seems to require that the phenomenal content of perceptual experience include observable properties—those properties we can perceptually ascribe as categorical. Yet this result stands in clear tension with 8. For if the Chalmers/Thompson argument goes through, no ordinary observable property—no shape, colour, size, and so on—will belong to the phenomenal content of perceptual experience. Given this tension, my Chapter 2 argument for PREDICATIVE ATOMISM undermines PHENOMENAL INTERNALISM (unless a proponent of the latter view can find some other way to resolve the tension).

Given the structure of the Chalmers/Thompson argument, a committed phenomenal internalist will be left few escape routes. I'll canvas three:

- I. Deny 1 by rejecting the possibility of displaced scenarios.
- II. Deny 2 by restricting PHENOMENAL INTERNALISM to a subset of observable properties

These two options require modifying the Chalmers/Thompson argument. Another option ('III') would be to keep their argument as it stands, and instead restrict the class of properties that subjects can perceptually ascribe as categorical. If the new class includes only properties immune to the Chalmers/Thompson argument, the phenomenal internalist would be home free.

I investigate these escape routes in more detail in §4.

4 No Prospects for PHENOMENAL INTERNALISM

This section surveys potential escape routes for proponents of PHENOMENAL INTERNALISM. Each falls under one of I–III. My aim here is modest. I seek to highlight the costs of I–III if one accepts both PHENOMENAL INTERNALISM and PREDICATIVE ATOMISM.

4.1 Edenic Sceptical Folly

Let's briefly recap the state of play. PREDICATIVE ATOMISM seems to require that phenomenal content include observable properties. But if Chalmers and Thompson are right, PHENOMENAL INTERNALISM forces us to banish these properties from phenomenal content (unless we wish to risk the sceptical threat of total illusion).

I want to begin with the third escape route—III—since it promises to preserve everything we might want: PHENOMENAL INTERNALISM, PREDICATIVE ATOMISM, the constraint imposed by PREDICATIVE ATOMISM on phenomenal content, and the Chalmers/Thompson argument. Implementing III requires isolating a set of properties (1) immune to the Chalmers/Thompson argument, but (2) whose perceptual ascription as categorical satisfies the demands of PREDICATIVE ATOMISM.

Chalmers (2006) introduces so-called ‘Edenic properties’ that satisfy (1). Here is how he introduces Edenic colour properties:

When I have a phenomenally red experience of an object, the object seems to be simply, primitively, *red*. The apparent redness does not seem to be a microphysical property, or a mental property, or a disposition, or an unspecified property that plays an appropriate causal role. Rather, it seems to be a simple qualitative property, with a distinctive sensuous nature. We might call this property perfect redness: the sort of property that might have been instantiated in Eden.

One might say: phenomenologically, it seems that visual experience presents the world to us as an Edenic world. Taking the phenomenology completely at face

value, visual experience presents a world where perfect redness and perfect blueness are instantiated on the surface of objects, as they were in Eden. These are simple intrinsic qualities whose nature we seem to grasp fully in perceptual experience. For the world to be *exactly* the way that my phenomenology seems to present it as being, the world would have to be an Edenic world in which these properties are instantiated.¹⁹

Edenic properties are meant to serve as the properties perceptual experience presents in virtue of its presentational phenomenal character. They play two roles. First, since Edenic properties are not presented merely as occupants of causal roles, their presentation captures the ‘categorical’ perceptual phenomenology that would otherwise fit poorly with a view like Fregean Representationalism which banishes ordinary observable properties from phenomenal content (e.g. we are inclined to say: ‘Red isn’t whichever colour plays such-and-such causal role—it is *that* colour’). Second, Chalmers uses Edenic properties to characterise the phenomenal character of perceptual experience. An experience with a ‘reddish’ phenomenal character, for example, will be one that presents a certain shade of Edenic red. But Chalmers argues, partly on scientific grounds, that Edenic properties are neither instantiated by ordinary objects nor play a substantial role in the natural causal order. So any attempt to attribute these properties to elements of the external world will inevitably lead to error, except when they are invoked as part of a characterisation of what it is like to undergo a perceptual experience.

Since Edenic properties figure in the phenomenal content of perceptual experience, no obvious barrier prevents us from perceptually ascribing them as categorical. Hence they appear to resolve the tension between PREDICATIVE ATOMISM and a version of PHENOMENAL INTERNALISM that accepts Edenic properties. But this appearance offers only false hope.

¹⁹ Chalmers (2006a, p. 66).

My argument for PREDICATIVE ATOMISM in Chapter 2 entails that perceptual ascription of Edenic properties as categorical generates an unacceptable form of scepticism. I argued for PREDICATIVE ATOMISM partly on the grounds that our capacity to ascribe special causal relations—those that involve the ordinary objects we perceive—rests upon our capacity to perceptually ascribe properties as categorical. For only the latter capacity enables us to isolate special causal relations in thought. However we cannot use Edenic properties to isolate actual special causal relations (as opposed to merely Edenic special causal relations), since these properties do not participate in the natural causal order. But without a capacity to ascribe these special causal relations, we cannot ascribe the causal relations that hold in virtue of unobservable properties. Thus if we can perceptually ascribe only Edenic properties as categorical, PREDICATIVE ATOMISM entails a bizarre sceptical picture on which we fail to ascribe non-Edenic unobservable properties. Yet proponents of PHENOMENAL INTERNALISM were driven to banish observable properties from phenomenal content in part because of their opposition to the possibility of such systematic error. So they cannot accept this sceptical position, let alone embrace it as a means to resist the threat of PREDICATIVE ATOMISM.

4.2 Inelastic Metaphysics

Options I and II both allow at least some ordinary observable properties into phenomenal content, yet retain some commitment to PHENOMENAL INTERNALISM. In principle, this general move would permit one to both perceptually ascribe non-Edenic observable properties as categorical—satisfying PREDICATIVE ATOMISM—and keep hold of the core of PHENOMENAL INTERNALISM. For example, the claim might be that we can perceptually ascribe shape properties (or relative distances between objects) as categorical, even if the Chalmers/Thompson argument precludes the perceptual ascription of size and colour properties as categorical.

Those who wish to make this move require a non-arbitrary way of drawing the line between those properties that ‘make it in’ to phenomenal content and those that do not. But I shall argue that

two factors that generate the tension between PREDICATIVE ATOMISM and PHENOMENAL INTERNALISM in the first place should make us quite pessimistic about the existence of such a non-arbitrary line within the class of observable properties. I'll discuss each factor in turn.

On the one hand, PHENOMENAL INTERNALISM puts pressure on us to exclude observable properties from phenomenal content. Chalmers and Thompson apply the argument I sketched earlier to a range of observable properties (spatial properties such as orientation, size, and shape; colour properties; and temporal properties like duration). If the argument works for any of these properties, it seemingly works for all of them. Any attempt to admit observable properties into phenomenal content must therefore explain why the Chalmers/Thompson argument applies only to the chosen subset of observable properties. And the prospects for a principled explanation of this sort look dim. Indeed, Chalmers has considered how the argument might extend to what at first glance look like potential exceptions: shape and relative size.²⁰ (For example, and put quite roughly, what look like squares in a displaced shape scenario might in fact be uniformly flattened relative to actual squares; or in a displaced scenario involving relative size I might see 1:4 ratios as 1:2 ratios, with everything else adjusted accordingly).

On the other hand, PREDICATIVE ATOMISM exerts pressure *against* attempts to exclude observable properties from phenomenal content. For any class of observable properties, if these properties cannot figure in phenomenal content, someone able to perceive only properties from that class would (given PREDICATIVE ATOMISM) be unable to ascribe unobservable properties. But there are hard questions about which observable properties (if any) are such that a subject able to perceive only those properties would be unable to ascribe unobservable properties. Furthermore, in light of my argument for the BRIDGE PREMISS in Chapter 2, subjects restricted to these properties would

²⁰ Cf. Chalmers (2012, pp. 328–32). For the shape case, he sketches a quite complicated scenario using resources from general relativity.

have to be unable to isolate special causal relations. While a failure to underwrite our grasp of causation might seem (somewhat) plausible for the small number of properties traditionally classified as ‘secondary qualities’—like colour or audible pitch—it seems much less plausible for properties traditionally classified as ‘primary qualities’ (shape, size, solidity, and the like).

Someone who wishes to save PHENOMENAL INTERNALISM therefore faces the following bind: if she accepts the Chalmers/Thompson argument, she risks two bad results. First, by excluding observable properties from phenomenal content, she might prevent phenomenal content from sustaining our grasp of special causal relations. Second, if we shift focus to the individual property categories themselves, our defender of PHENOMENAL INTERNALISM risks saying something about primary qualities that plausibly applies only to secondary qualities like colour (namely that grasp of these properties does not permit a subject to grasp causal relations).

But suppose a fan of PHENOMENAL INTERNALISM wishes to brave these difficulties. In order to admit some ordinary observable properties into phenomenal content yet retain commitment to some version of PHENOMENAL INTERNALISM, such a theorist has two options. We’ve already listed them as (I) and (II): (I) deny that displaced scenarios involving these properties are at all possible (that is, deny 1 from the Chalmers/Thompson argument); (II) deny that a subject in a displaced scenario involving these properties has perceptual experiences type-identical to our own (that is, deny 2 from the Chalmers/Thompson argument).²¹ Given the difficulties already raised, I’m not sanguine about either of these options. Yet I cannot discount them altogether—some enterprising soul might ultimately find a way around the difficulties I’ve derived from PREDICATIVE ATOMISM. The ball now lies in the court of PHENOMENAL INTERNALISM.

²¹ Prosser (2011), for example, pursues a version of (I). He argues that the phenomenal content of perceptual experience includes affordance properties. These are properties individuated partly in terms of a subject’s practical or agentive capacities: an object might for example be *graspable*, or *twistable*, or *near*. He claims these properties remain invariant through displacement.

5 Prospects for PHENOMENAL EXTERNALISM

Ours has been a winding road. I first argued that in securing a central role for our capacity to perceptually ascribe observable properties as categorical, PREDICATIVE ATOMISM plausibly ensures that perceptual experience must include observable properties within phenomenal content.

Yet this result quickly came into conflict with PHENOMENAL INTERNALISM, the view on which what it is like for a subject to undergo a perceptual experience supervenes upon the subject's intrinsic properties (and so does not depend upon which objects or properties are in the subject's environment, or upon her relationship to these entities). We saw an argument from David Chalmers and Brad Thompson which seemed to establish that PHENOMENAL INTERNALISM requires banishing ordinary observable properties from the phenomenal content of perceptual experience.

I canvassed three ways for proponents of PHENOMENAL INTERNALISM to escape the tension between their view and PREDICATIVE ATOMISM. I showed that one of these escape routes simply fails (the appeal to so-called 'Edenic' properties), and that the other two face serious difficulties. Given these difficulties, the combination of PHENOMENAL INTERNALISM and PREDICATIVE ATOMISM seems almost untenable.

Given my argument for PREDICATIVE ATOMISM in Chapter 2, I suggest we jettison PHENOMENAL INTERNALISM. We should instead adopt some type of PHENOMENAL EXTERNALISM on which facts about perceptual presentation—and thus about presentational phenomenal character—depend in part upon which properties are actually in the environment. Since what it is like for a subject will depend in part on the actual properties her experience presents, two experiences may differ in presentational phenomenal character yet remain indiscriminable from a subject's point of view. A natural way to embrace this result would be to adopt direct realism. The view entails PHENOMENAL EXTERNALISM: in veridical cases the objects and properties to which an experience relates a subject combine with relevant features of the subject's perceptual perspective to determine what it is like for the subject. However PHENOMENAL EXTERNALISM does not

straightforwardly entail direct realism (for one thing, the version of PHENOMENAL EXTERNALISM licensed by the argument of this chapter remains silent about whether presentational phenomenal character also depends upon *objects*).²²

PHENOMENAL EXTERNALISM classifies total illusions as impossible on the grounds of incoherence. If the phenomenal character of a perceptual experience depends in part upon which non-Edenic observable properties it presents, the typing of a subject's perceptual experiences presupposes facts about what sort of scenario the subject occupies. For example, a specification of what it is like for me at the moment would appeal to my laptop's shape, my lamp's colour, and so on. Total illusions are thus incoherent because a systematic mismatch between a subject's perceptual experience and the world cannot be consistently maintained. We cannot both hold fixed what it is like for a subject and insist that she actually lives in a displaced scenario. An adequate specification of what it is like for a subject will presuppose that her environment contains particular observable properties. So to insist that she lives in a displaced scenario, and thus that her environment does not actually include those properties, would be to assert an inconsistent description of her situation (one according to which she both is, and isn't, in a non-displaced scenario). In contrast, PHENOMENAL INTERNALISM assumes that the phenomenal character of a perceptual experience carries no intrinsic commitment to one scenario over another.

6 Looking Back to Look Forward

I wish to end on a synoptic note that connects the argument of this chapter with the work done in other chapters. In this chapter I've brought out a tension between PREDICATIVE ATOMISM (the flagship claim of Chapter 2) and PHENOMENAL INTERNALISM. Our best escape is to reject PHENOMENAL INTERNALISM, and instead move towards a sort of PHENOMENAL EXTERNALISM—a

²² Johnston (2004), for example, adopts PHENOMENAL EXTERNALISM, but stops short of direct realism. Others (such as Martin 2004) argue that any such intermediate position is untenable.

view on which the phenomenal character of perceptual experience is partially fixed by which properties are present in the perceived environment. This sort of PHENOMENAL EXTERNALISM saves us from total illusion, and does not require the perceptual alienation from observable properties that led PHENOMENAL INTERNALISM to run afoul of PREDICATIVE ATOMISM.

This result dovetails nicely with the arguments of Chapter 1. I argued there that a right account of visual experience ought to begin with claims about which objects and properties an experience puts a subject into perceptual contact with. Yet a viable notion of perceptual contact with external world objects and their properties becomes difficult to maintain when saddled with PHENOMENAL INTERNALISM. Whether a subject stands in perceptual contact with a given object or property isn't something that she can reliably recognise from her own point of view. Thus the argument against PHENOMENAL INTERNALISM developed in this chapter—an argument that owes its foundation to the argument for PREDICATIVE ATOMISM in Chapter 2—removes a potential obstacle facing accounts of visual experience developed along the lines suggested at the end of Chapter 1.

Finally, a brief look ahead: in the next chapter I'll raise a challenge to the suggestion that we genuinely think about properties at all, and show how the view I've developed over the last three chapters goes some way to address the challenge.

Chapter 4

The Limits of Reference

This chapter resurfaces a neglected challenge originally articulated by Michael Dummett (1973/1981) to the claim that we genuinely think and talk about properties. Though Dummett raised it as an internal challenge within Frege's framework, I show that the challenge remains standing even if we jettison central Fregean commitments.

Dummett's target was Frege's insistence that so-called 'incomplete' expressions refer.¹ For example, given the sentence 'John is heavy' (a 'complete' expression), Frege takes both the name 'John' (another complete expression) and the incomplete expression ' ξ is heavy' (derived from 'John is heavy' by dropping 'John') to refer to extra-linguistic entities. Dummett's challenge rests on an apparent tension that emerges when we ask how two central Fregean claims should be understood:

1. All (non-degenerate) significant expressions possess both sense and reference (where 'significant' expressions are just those that could be understood).
2. The senses of incomplete expressions are themselves incomplete (or 'unsaturated').

The bare bones of Frege's distinction between sense and reference are familiar. The referent of an expression is its contribution to the truth or falsity of sentences within which it occurs. The sense of an expression—what a speaker grasps when she understands the expression—is a mode of presentation of its referent. But we'll see that when we consider the relationship between sense and reference in more detail, we find that Frege is also committed to the explanatory priority of the level of sense over the level of reference:

3. An attribution of reference to a class of expressions is justified only if it is required by a right account of their senses.

¹ Cf. Frege (1979, pp. 118–19).

2 and 3 entail that an attribution of reference to incomplete expressions will be justified only if it is required by a right account of incomplete senses. Yet it turns out that the puzzle that motivates the adoption of incomplete senses in the first place (what I later dub the ‘Unity of Understanding Problem’) also supports 4:

4. An account of the incomplete sense of an incomplete expression does not require the attribution of reference to the expression.

2–4 entail 5:

5. The attribution of reference to incomplete expressions is unjustified. [From 2–4]

Yet 5 and 1 conflict; we have our tension.

Dummett’s formulation of the challenge presupposed a verificationist account of the senses of singular terms and incomplete expressions. This has doubtless contributed to the general failure to recognise the significance of his challenge. I therefore use §1 to show how, even if we ditch Dummett’s verificationist account of Fregean sense, the challenge still arises within Frege’s system. I then argue in §2 that the challenge remains even when we abandon many of Frege’s central commitments. As a result, what begins as a problem within Frege’s framework becomes a general challenge to contemporary attributions of reference to predicative expressions, as well as to parallel attributions at the level of thought. I survey a range of potential responses to the challenge in §3. I develop my own solution to the challenge in the conclusion of the dissertation.

1 The Challenge in Frege’s Framework

In this section I show how Frege’s framework commits him to 1–4. In §1.1 I say just enough about Frege’s distinction between sense and reference to motivate 1 and 3. In §1.2 I explain Frege’s introduction of incomplete senses (which commits him to 2). Finally, in §1.3 I argue that a Fregean account of incomplete senses does not require an attribution of reference to incomplete expressions—but this is just 4.

1.1 Reference and Understanding

In this subsection I trace how Frege's distinction between sense and reference commits him to 1 (namely: all non-degenerate significant expressions possess both sense and reference) and to 3 (namely: the attribution of reference to a class of expressions is justified only if it is required by an account of the senses of these expressions). I begin with a relatively bare bones characterisation of his sense/reference distinction.

1.1.1 Frege on Sense and Reference

Frege introduced the notion of reference in part to explain how the truth or falsity of a sentence depends upon its constituent expressions. He proposed that a sentence's constituent expressions have referents, and the truth or falsity of the whole sentence is determined by the referents of its constituent expressions by functional application. Hence if a sentence has a truth-value, all of its constituent expressions have a referent. A sentence that contains an expression without a referent is thus a degenerate case: such a sentence has no truth-value.

This link between reference and truth generates Frege's criterion for co-reference. Less precisely: two expressions co-refer iff they make the same contribution to the truth-values of sentences within which they occur. More precisely:

Co-Reference: Two expressions e_1 and e_2 possess the same referent iff wherever ' $\dots \dots$ ' is a non-oblique context, ' $\dots e_1 \dots$ ' is true iff ' $\dots e_2 \dots$ ' is true.²

An oblique context is one in which expressions do not possess their usual referent (for instance, quoted expressions refer to themselves).

² Cf. Frege (SR, p. 59). All page references to Frege's papers 'On Sense and Reference', 'On Concept and Object' and 'Function and Concept' ('SR', 'CO', and 'FC' respectively) are from Black and Geach (1952); references to *Grundgesetze der Arithmetik Vol. I* (hereafter 'GgI') are from Beaney (1997). Also, a brief remark about notation: the quotation marks are corner-quotes. They should be read as follows: ' α is Φ ' is an atomic sentence' says the left-to-right concatenation of a name denoted by ' α ', the copula 'is', and a predicate denoted by ' Φ ' is an atomic sentence.'

Frege introduced the notion of sense to solve a puzzle about speaker understanding—*Frege's Puzzle*. His main version of the puzzle involved singular terms. Take two co-referential proper names (say ‘Cicero’ and ‘Tully’). Now compare the situation of two thinkers, Thomas and Nicole, both of whom count as competent users of ‘Cicero’ and ‘Tully’. Thomas believes that Cicero is wise but also disbelieves that Cicero is wise. In contrast, Nicole believes that Cicero is wise but disbelieves that Tully is wise. Intuitively, Nicole has committed a less egregious error than Thomas. She merely has two beliefs that cannot both be true at once (that Cicero is wise and that Tully is not wise), while Thomas seems guilty of a form of rational incoherence. But we cannot explain this difference between Thomas and Nicole in terms of reference alone: ‘Cicero is wise’ and ‘Tully is wise’ contain only co-referential constituent expressions.

Though the most well known Frege Puzzles concern singular terms, an exactly parallel puzzle arises for the case of incomplete expressions. Take two predicates which we might intuitively regard as standing for the same property. For example, ‘is a natural number less than 2’ and ‘is a natural number less than the cube root of 8’. Just as we can imagine two thinkers both of whom count as competent users of two names, but who exhibit a contrast in the way Thomas and Nicole do, we can have two thinkers one of whom believes 1 is a natural number less than 2, and also disbelieves that 1 is a natural number less than 2; and another who believes that 1 is a natural number less than 2, but also disbelieves that 1 is a natural number less than the cube root of 8. One of these thinkers is guilty merely of having two beliefs that cannot be true at once (like Nicole), while the other seems guilty of a form of rational incoherence (like Thomas).

Frege’s notion of sense is therefore meant to explain how thinkers could adopt conflicting attitudes to what is expressed by pairs of sentences which differ one from the other only in the substitution of a co-referential constituent expression. So we get a criterion for sameness of sense:

Sameness of Sense: Two expressions e_1 and e_2 possess the same sense iff for a subject S who understands e_1 and e_2 , S cannot both accept $\ulcorner \dots e_1 \dots \urcorner$ and reject $\ulcorner \dots e_2 \dots \urcorner$ without loss of rational coherence.³

Frege says that the sense of a sentence is a ‘Thought’, while the senses of the sentence’s sub-sentential expressions are the contributions of these expressions to the Thought.⁴ A thinker can accept ‘Cicero is wise’ but reject ‘Tully is wise’ without rational incoherence because these sentences express distinct Thoughts.

My relatively sparse statement of Frege’s distinction between sense and reference suffices to show that Frege was committed to 1, the claim that every non-degenerate significant expression possesses both a referent and a sense. If a significant expression were to lack a referent, sentences containing it would lack truth-values; and if it were to lack a sense, sentences containing the expression would fail to express Thoughts. Sentences that lack a truth-value or fail to express a Thought are degenerate. So a significant expression that lacks a referent or a sense is also degenerate.

1.1.2 Two Aspects of Reference

With the bare bones of Frege’s distinction between sense and reference on the table, we’re in a position to introduce a distinction that Dummett recognised between two aspects of the Fregean notion of reference.

The first aspect is already visible in the sparse statement of the distinction between sense and reference given in the previous subsection. The referent of an expression is its contribution to the determination of the truth-values of sentences within which it occurs—what Dummett calls its

³ Frege (SR, pp. 56, 62). There is a debate about how to refine this criterion, and so the notion of sense, but nothing I say here depends on which side we choose. Apart from Dummett (1973/1981; 1981), criteria for sameness of sense are offered by Evans (1982), Rumfitt (1994), and Dickie and Rattan (2010).

⁴ I’ve adopted the widespread convention of using ‘Thought’ rather than ‘thought’ for the sense of a sentence. The convention will help me avoid potential confusion when I turn to mental representation in §2.

‘semantic value’.⁵ It is this aspect of reference that underwrites Frege’s claim that if a sentence has a truth-value, all of its constituent expressions have a referent. For if an expression lacks a semantic value, it cannot contribute to the determination of the truth-values of sentences within which it occurs.

Yet Dummett’s challenge does not concern whether incomplete expressions possess reference in this thin sense. The attribution of reference *as semantic value* to incomplete expressions falls right out of Frege’s link between reference and truth. Incomplete expressions are significant, and figure in true sentences, so they must possess a semantic value.

The challenge instead concerns the second aspect of Frege’s notion of reference, which I shall now explain. Dummett observes that Frege explained all reference in terms of an analogy to the prototypical relation between a name and its bearer.⁶ The name/bearer relation ties the purely formal theory of semantic value to the world. While the primary constraint on a pure theory of semantic value is the need to assign correct truth-values to sentences, this constraint does not force expressions to stand for denizens of the external world. The identification of the referent of a name with its bearer tells us what a pure theory of semantic value cannot, namely that there are specific objects such that a name’s contribution to the truth values of sentence is to be explained by the fact that the name stands for that specific object rather than for something else.

The significance of Frege’s name/bearer prototype for the theory of sense emerges from his doctrine that sense determines reference. To specify what kind of thing an expression stands for, as Frege does when he says that a name stands for its bearer, places constraints on a right account of

⁵ Dummett actually shifts his terminology over time. When introducing his challenge in Dummett (1973/1981, pp. 190–1), he speaks of an expression’s ‘semantic role’ rather than its ‘semantic value’. He later switches to talk of ‘semantic value’ in Dummett (1991a, pp. 30–2).

⁶ As Furth (1968) points out, Frege’s system cannot admit just a single relation of reference. Hence any ‘referential’ relation between incomplete expressions and functions can only be analogous to that between names and their bearers. More recently, several philosophers have argued for a sharp distinction between the semantic relationships between names and their bearers, on the one hand, and predicative expressions and their semantic values, on the other hand. See Liebesman (2015) and Rieppel (forthcoming).

that expression's sense.⁷ For example, the sense of a name must be a mode of presentation of its bearer (rather than some other object, like the singleton set containing the name's bearer).⁸ Given this framework, anything to do with the attribution of reference to a class of expressions that goes beyond the mere attribution of semantic value must come from a right account of the senses of these expressions (and thus from an account of what it is to understand them). But this is just 3—the claim that the attribution of reference to a class of expressions is justified only if this attribution is required by a right account of the senses of these expressions.⁹

What is it for an attribution of reference to be required by a theory of sense? Dummett's model was an account on which a speaker understands a singular term only if she grasps a criterion of identity—a criterion that tells a speaker what it would take for a given object to be identical to the bearer of the singular term.¹⁰ Hence the core of Dummett's challenge to Frege for the case of incomplete expressions is the claim that understanding of an incomplete expression does not share this central feature of understanding of singular terms. That is, according to Dummett a speaker who understands an incomplete expression need not—and indeed cannot—be credited with grasp of a criterion of identity (or anything sufficiently like it) for the referent of the expression.

Dummett unfortunately gave a verificationist account of a speaker's grasp of a criterion of identity: put roughly, to possess a criterion of identity consists in an ability to decide, for any given

⁷ Dummett (1981, p. 249).

⁸ A natural way to implement the thesis that names are predicates—a position currently enjoying a revival—is to say that the semantic value of a name is the set containing its bearer (or bearers). See Quine (1960, p. 163), Burge (1973), and Fara (2015).

⁹ As Dummett (1991a, p. 24) points out, there is little question that these two distinct roles for the notion of reference exist in Frege's mature philosophy of language. See Frege (SR, p. 64) for the claim that a name's referent is its contribution to the determination of the truth-value of sentences within which the expression occurs, and Frege (SR, 57) for the claim that the sense of a name is a mode of presentation of the name's bearer. Frege is distinctive in part for his identification (rather than separation) of what satisfies these two roles.

¹⁰ In appealing to criteria of identity Dummett follows Frege (1984, §62).

object, whether it is the bearer of the name.¹¹ Given the widespread resistance to Dummett's verificationism, its lurking presence has doubtless prevented many from recognising the strength of Dummett's challenge to Frege's attribution of reference to incomplete expressions.

However verificationism isn't required to get the challenge off the ground. Dummett's appeal to criteria of identity suffices to forge a link between the attribution of reference and a theory of sense, even without his own verificationist gloss. For the attribution of reference to a singular term determines *which* criterion of identity a speaker must possess to understand the expression. Verificationism merely provides a mechanism that explains how speakers might satisfy one end of this link between reference and sense. In principle, one could accept Dummett's appeal to criteria of identity yet offer a non-verificationist account of what it is to possess a criterion of identity. Non-verificationist neo-Fregeans often develop accounts of just this sort. They defend versions of the idea (rejected by Dummett) that possession of a criterion of identity consists in knowledge of what it would be for an object to be the bearer of a name.¹²

This distinction between two aspects of Frege's notion of reference is a central pillar of Dummett's challenge. I therefore use the remainder of this subsection to consolidate it.

Few insist that our intuitive conception of aboutness—a conception to which Frege's notion of reference is answerable—is exhausted by a conception of reference as semantic value. Yet some may argue that the name/bearer prototype introduces implausibly strong requirements on the

¹¹ Dummett (1973/1981, pp. 230, 488). Here is Dummett's refined formulation on p. 230: 'to grasp the sense of a name is to be able to recognise, when one is presented with it, whatever counts as establishing conclusively that a given object is the bearer of the name'. Of course, Frege wasn't a verificationist (as Dummett acknowledges). Cf. Frege (1984, §62).

¹² Gareth Evans (1982, Ch. 4) gives a non-verificationist account of this sort. Following Strawson (1959, Ch. 1), Evans holds that knowledge of what it would be for an object to be the bearer of a name requires discriminating knowledge: 'the subject must have a capacity to distinguish the object of his judgment from all other things.' (p. 89). (An early model of discriminating knowledge crops up at the end of Plato's *Theaetetus*; see Burnyeat 1990 for discussion.) On Evans's account—which in important respects parallels Dummett's account of names for abstract objects (cf. Dummett 1973/1981, pp. 489, 498–9)—discriminating knowledge rests on possession of a 'fundamental Idea' of an object. To possess a fundamental Idea of an object is to think of the object as possessing a 'fundamental ground of difference' appropriate to things of its sort, where this ground is a feature that differentiates the object from all others. Evans argues that the fundamental ground of difference for ordinary objects is an object's location and kind.

attribution of reference to expressions outside the class of singular terms. They might instead insist that our intuitive conception of aboutness requires something weaker: merely the association of an expression with some extra-linguistic entity out in the world (a name with its bearer; a sentence with its truth-value; etc.).

One way to establish this weaker requirement would be to follow Frege's explanation of quantification. Frege originally introduced the notion of reference to help explain the truth conditions of quantification-involving sentences such as 'Something is tall'.¹³ Quantifiers divide into at least two categories. First order quantifiers (such as the 'something' in 'Something is tall') quantify into the sentence position occupied by singular terms, since (to follow Frege) we introduce these quantifiers by first dropping one or more singular terms from a sentence: for example, starting with the sentence 'that is tall', we get ' ξ is tall'; the gap in this expression can be filled either with a name (to get 'Jack is tall') or filled with a quantifier to get 'something is tall'. Second order quantifiers, in contrast, quantify into the sentence position occupied by predicates ('Dawn and Ed are something I am not—namely hungry').

Frege explains quantification in terms of the referential role of the expressions that occupy the sentence positions the quantifiers quantify into. For example, a sentence such as 'something is tall' will be true just in case some object—a potential bearer of a singular term—is tall. Similarly, we could say that the sentence 'Dawn and Ed are something I am not' will be true just in case some property—a potential referent of a predicate—is possessed by Dawn and Ed but not by me. So if sentences that involve quantification into predicate position are intelligible, and if such

¹³ Cf. Frege (1879, §11).

quantification should be explained by appeal to a range of extra-linguistic entities (properties, for example), we have reason to assign incomplete expressions extra-linguistic ‘referents’.¹⁴

Yet this kind of appeal to quantification—or to any other resource that remains outside our theory of sense—cannot capture our intuitive conception of aboutness (a conception against which any adequate Fregean account of reference must be measured). The referent of a name (the base case for both Frege and Dummett) cannot be any old object whose association with the name permits us to generate appropriate truth conditions for sentences involving first order quantification. Within Frege’s framework, a name’s referent must be the object presented to a subject when she grasps its sense (or, more precisely, the object whose criterion of identity a subject must grasp to understand the name). Dummett’s challenge puts pressure on Frege’s extension of this rich notion of reference to incomplete expressions—a notion grounded in the dual role of reference as semantic value and as modelled by the name/bearer prototype. No appeal to second order quantification can alone settle this challenge. We must instead show that a right account of incomplete senses requires the attribution of reference to incomplete expressions.

1.2 Incomplete Senses

I’ve spelled out how Frege’s distinction between sense and reference commits him to elements 1 and 3 of the challenge presented at the beginning of the chapter. This subsection introduces and motivates Frege’s commitment to 2: the claim that the sense of an incomplete expression is itself incomplete (or ‘unsaturated’). I begin with the composition of Fregean Thoughts in §1.2.1: Dummett offers a verificationist account, but I (once again) argue that one can retain the essential components of Dummett’s account without verificationism. With an account of the composition of Fregean Thoughts on the table, I motivate and explain Frege’s complete/incomplete distinction. I

¹⁴ Those who accept this sort of link between quantification and reference include Frege (1879), Russell (1918/1956), Wittgenstein (1921), Dummett (1973/1981), and Evans (1982). For recent relevant discussion of second order quantification, see Rayo and Yablo (2001), Wright (2007), and Rieppel (forthcoming).

begin with the distinction at the level of language—Frege’s ‘home case’ (§1.2.2). I then show how Frege extends this distinction first to the level of reference (§1.2.3), and then to the level of sense (§1.2.4). Frege’s extension of the distinction to the level of sense (which commits him to 2) depends in part upon the account of Thought composition sketched in §1.2.1.

1.2.1 Composition of Fregean Thoughts

A key difference between sense and reference lies in the relationship between Thoughts and truth-values, on the one hand, and the sense and reference of sub-sentential expressions, on the other. The referents of sub-sentential expressions functionally determine the truth-value of complete sentences: the truth-value does not contain these referents as parts. In contrast, Thoughts are *composed* of senses.¹⁵ The distinction between sense and reference therefore also represents a distinction between the kinds of structure present at each level.

I follow Dummett’s account of what it is for senses to compose Thoughts. Let us say that a speaker understands a token sentence iff she ‘grasps’ the Thought it expresses.¹⁶ Dummett explains the composition of Thoughts in terms of what it is to grasp them. To grasp the Thought expressed by a sentence is to know the conditions under which the sentence is true. Hence because the truth-value of a sentence is determined by the referents of its constituent expressions, to grasp the Thought expressed is to know how the referents of the sub-sentential expressions, each presented by a corresponding sense, must combine for the sentence to be true.¹⁷ So to grasp the sense of a sub-sentential expression is to associate the expression with a specific kind of contribution to the truth conditions of sentences within which it occurs.

¹⁵ These claims can be found at Frege (SR, p. 65; GgI, §32). For discussion, see Heck and May (2011).

¹⁶ Perry (1977) raises difficulties for Frege’s account of the sense of indexicals such as ‘I’ or ‘here’, since tokens of an indexical often make distinct contributions at the levels of sense and reference. See Evans (1981) for a powerful defence of Frege.

¹⁷ See Dummett (1973/1981, p. 152).

As with the name/bearer prototype, Dummett's account of the composition of Thoughts begins with an attractive abstract claim—in this case the claim that we can explain the composition of a Thought in terms of the capacities we exercise when we grasp the Thought—and then proceeds to give this claim a verificationist gloss. As we've seen, grasp of a Thought consists in grasp of the conditions under which it is true. Yet Dummett thinks that to grasp the Thought expressed by 'that is round', for example, a subject must grasp a *procedure* for determining whether the sentence is true or false. And this procedure will combine two more primitive procedures: one for identifying the bearer of 'that', and another for determining whether 'round' applies to the bearer of 'that'. A subject who grasps these procedures counts as grasping the senses of 'that' and 'round', respectively.

Non-verificationists will look askance at Dummett's assimilation of our grasp of sense to grasp of a mere procedure. Yet as was the case with the name/bearer prototype, retaining Dummett's valuable non-verificationist insight (namely that the composition of Thoughts should be explained in terms of the capacities exercised in grasping these Thoughts) does not require his verificationism. Grasp of the Thought expressed by 'that is square', for example, could instead be constituted by joint exercise of two capacities: a capacity to think about the bearer of 'that' and a capacity to think ξ is *square* of some arbitrary member of the domain. In place of Dummett's procedures, thinking about a particular object simply requires grasp of a certain *ways of thinking* about the object.¹⁸ Given the availability of this sort of non-verificationist account of Thought composition—an account that preserves Dummett's essential insight—we've removed yet another vestigial element of verificationism from Dummett's challenge to the attribution of reference to incomplete expressions.

¹⁸ Evans (1982, Ch. 1) pioneered this non-verificationist explanation of Frege's notion of sense. Indeed, Evans also argues that Dummett's verificationism isn't essential to his account of the composition of Thoughts (see pp. 16-17 of Evans 1982, especially the long discussion of Dummett in n. 17). Evans's account of sense in terms of 'ways of thinking' remains quite influential. Recent adherents include Dickie and Rattan (2010), Stanley (2011), and Recanati (2012).

Now let us push these issues about verificationism back into the closet, since we still must complete the last bit of Dummett's account of Thought composition. The relationship between Frege's theory of reference and his theory of sense also reflects the dual role of sense. First, an expression's sense is a mode of presentation of its referent (hence its referent is the entity its sense presents). Second, the notion of sense plays a role in the structure of Thoughts. Thoughts are composed of senses. So a Thought is composed of modes of presentation of the referents of the constituents of the sentence that expresses it. In the case of language, to specify the contribution of an expression's sense to Thoughts is to explain how we must use the expression in order to form sentences that express Thoughts.¹⁹

1.2.2 The Complete/Incomplete Distinction Within Language

While Frege eventually extends his complete/incomplete distinction to the levels of sense and reference, his initial application of the distinction is to language. Here's an example to see how Frege thinks the distinction works at the level of language. Begin with the sentence 'Gertrude likes Bertie'. Starting with this sentence, we can form the incomplete expression ' ξ likes ζ ' by dropping the two names (where ' ξ ' and ' ζ ' signify argument-places).²⁰ Such an expression is 'incomplete' because of its derivative status: it is extracted from the original sentence by dropping the two names. And the way an incomplete expression is obtained from the initial sentence will constrain what counts as a grammatical sentence built out of the expression.

Singular terms and sentences are 'complete' because they do not place the same constraints on the construction of sentences as incomplete expressions do. Given a singular term, there is no determinate fact about how many argument-places an incomplete expression must have to combine with the singular term to form a grammatical sentence (e.g. 'Gertrude' could be used in combination

¹⁹ For this same point, see Dummett (1973/1981, pp. 146–7).

²⁰ Cf. Frege (1879, §9). For discussion of Frege's distinction between complete and incomplete expressions, see Dummett (1973/1981, Ch. 2); Wright (1983, pp. 7–8); Rumfitt (1994).

with either ‘ ξ likes ξ ’ or ‘ ξ likes Bertie’ to form a sentence). In contrast, given an incomplete expression, there is a determinate fact about how many singular terms or variables, and what kinds of singular terms or variables, we can combine with the incomplete expression to form grammatical sentences.

1.2.3 How Frege Extends Incompleteness to Reference

Frege’s distinction between complete and incomplete expressions forms the basis of an asymmetric treatment of the sense and reference of sub-sentential expressions. At the level of reference we get an exhaustive division between two ontological categories: ‘Objects’ are what complete expressions (singular terms and sentences) refer to; ‘functions’ (mind-independent mappings from arguments to values) are what incomplete expressions, of which predicates are the base category, refer to. In the special case of predicates—Incomplete expressions containing only a single argument place—their referents are ‘concepts’. These are mind-independent functions from objects to truth-values (relations, in contrast, are functions from pairs of objects to truth-values).²¹

The distinctions between complete and incomplete expressions, on the one hand, and between complete and incomplete entities, on the other, rest on a common explanatory structure. Incomplete expressions are individuated in terms of their role in sentences, while incomplete entities (functions) are individuated in terms of their role (as potential referents) in the determination of truth-values (the referents of sentences). This explanatory structure accords priority to the ‘complete’ category: the ‘incomplete’ category is explained in terms of its relationship to members of the complete category.

²¹ For the claims about reference, compare Frege (1984, §57 and §60; FC, pp. 24–5; 1979, pp. 172–3). For the extension of the complete/incomplete distinction to the level of sense, see Frege (CO, pp. 54–5).

1.2.4 How Frege Extends Incompleteness to Sense

Frege also adopted a complete/incomplete distinction at the level of sense. The extension of the complete/incomplete distinction to the level of sense requires the retention of this structure of explanatory priority. But instead of the composition of sentences or the determination of truth-values, for the case of senses Frege's target is the composition of Thoughts.

Keeping the parallel to the case of reference, ‘complete’ (or ‘saturated’) senses are those expressed by complete expressions (and so are either senses of singular terms or Thoughts expressed by sentences); ‘incomplete’ (or ‘unsaturated’) senses are those expressed by incomplete expressions. Keeping the parallel to language, incomplete senses enjoy a derivative status: given a sentence that expresses a Thought (itself a ‘complete’ sense), we can ‘drop’ from the Thought the ‘complete’ sense of one of the sentence’s singular terms. What remains is an incomplete sense—specifically the incomplete sense expressed by the incomplete expression formed from the original sentence by dropping the relevant singular term.

Aside from these formal parallels, the complete/incomplete distinction at the level of sense plays a central role in giving Frege his solution to a problem that arises when we consider what is involved in understanding sentences (or equivalently, what it is to grasp Thoughts). Here is the problem:

UNITY OF UNDERSTANDING PROBLEM: The Thought *a* is *F* is composed of modes of presentation of **a** and **F**. A thinker who grasps *a* is *F* counts as thinking that **a** falls under **F**. The alternative is that she counts as thinking of a mere list of entities (the list consisting of **a** and **F**). In virtue of what does she count as thinking that **a** falls under **F**?

I take Frege to have this problem in mind in his discussion of incomplete senses in the penultimate paragraph of ‘On Concept and Object’ when he says:

For not all the parts of a thought can be complete; at least one must be ‘unsaturated,’ or predicative; otherwise they would not hold together. For example,

the sense of the phrase ‘the number 2’ does not hold together with that of the expression ‘the concept prime number’ without a link. We apply such a link in the sentence ‘the number 2 falls under the concept prime number’; it is contained in the words ‘fall under,’ which need to be completed in two ways—by a subject and an accusative; and only because their sense is thus ‘unsaturated’ are they capable of serving as a link.²²

It is tempting to read this passage as addressing a purely metaphysical question about what it is for the constituents of a Thought to ‘hold together’. But such a reading ignores the role of Thoughts in an account of speaker understanding: to understand an expression is to grasp its sense. We saw in §1.2.1 how a Thought is essentially a reification of the capacities a subject exercises when grasping it. So any special question of the unity of the Thought is really a question of the unity of understanding.²³

The UNITY OF UNDERSTANDING PROBLEM demands that the role of $\zeta \text{ is } F$ be explained in terms of what it is to grasp $a \text{ is } F$ and a . To grasp the Thought $a \text{ is } F$ is to know the conditions under which it is true. $a \text{ is } F$ is true iff the incomplete sense $\zeta \text{ is } F$ is true of the object that a stands for. This construal of the truth conditions of $a \text{ is } F$ requires that $\zeta \text{ is } F$ be regarded as having a semantic value—its contribution to determining truth or falsity of Thoughts in which it occurs—but does not give us any reason to think $\zeta \text{ is } F$ has a referent in the richer name/bearer sense. (Note: while I speak loosely here and elsewhere of the referent of a sense, this is merely shorthand for ‘the referent of the expression to which the sense belongs.’) The contribution of a to the truth conditions of the thought is to introduce an object for the thought to be about, and so a thinker who grasps a must know what it would be for an object to be the referent of a .

²² Frege (CO, p. 54).

²³ I am not the first to have seen this connection in Frege. Rumfitt (1994, p. 611) argues that ‘For a proper understanding of those propositions that contain ‘contradicts’ as a part, it is not enough to know that the word expresses the relational property of contradicting. Rather, one needs to know that that a proposition formed by concatenating n with ‘contradicts’ with m means that n’s bearer contradicts m’s.’

$\zeta \text{ is } F$ cannot do the same work as a (namely merely to introduce an entity) on pain of rendering the UNITY OF UNDERSTANDING PROBLEM unsolvable. To grasp a is F would amount to no more than the identification of a pair of entities—a mere list. What it is to grasp $\zeta \text{ is } F$ must therefore reflect the role of $\zeta \text{ is } F$ as something which is true or false of objects. Hence a thinker who grasps $\zeta \text{ is } F$ must be in a position to know what it would be for $\zeta \text{ is } F$ to be true of a given object.²⁴ Combined with knowledge of what it would be for an object to be the referent of a , the knowledge of a thinker who grasps $\zeta \text{ is } F$ constitutes knowledge of what it would be for $\zeta \text{ is } F$ to be true of the referent of a . But this is just knowledge of what it would be for a is F to be true.

The complete/incomplete distinction at the level of sense can be explained in terms of what it is to grasp a and $\zeta \text{ is } F$. A thinker's grasp of a must, if we are to answer the UNITY OF UNDERSTANDING PROBLEM, reflect the fact that it has to be combined with grasp of an incomplete sense in order to constitute grasp of a Thought. But grasp of a does not need to impose substantial constraints on the nature of the incomplete senses that a can be combined with. In contrast, an incomplete sense places constraints on the number and type of complete senses it can combine with to form Thoughts (remember that a similar set of constraints applies at the level of language). So grasp of $\zeta \text{ is } F$ must involve knowledge that reflects these more substantial constraints on the kinds of Thoughts $\zeta \text{ is } F$ can be used to form. But such knowledge will ultimately be explained in terms of our grasp of what it is for $\zeta \text{ is } F$ to be true or false of objects introduced by the rest of the Thought.

²⁴ Dummett (1973/1981, pp. 223–34), Evans (1982, p. 109), and Peacocke (1992, pp. 43–4) all endorse a version of this treatment of our grasp of $\zeta \text{ is } F$, though they provide distinct accounts of our knowledge of what it would be for $\zeta \text{ is } F$ to be true of an arbitrary object. (For a non-Dummettian defender, see Reimer 2002; and for scepticism about the view, see Travis 1994.) Dummett's verificationist account leads to anti-realism, while Evans and Peacocke strive to articulate accounts friendly to realism.

1.3 No Need for Predicative Reference

I've traced how Frege's framework commits him to the first three components of Dummett's challenge—a challenge that so far does not depend upon verificationism. It remains to show that Frege's framework commits him to the final component, namely the claim that a right account of incomplete senses does not require the attribution of reference to their associated incomplete expressions. I have three arguments that together converge on this claim.

1.3.1 An Appeal to the History of Metaphysics

For a *prima facie* reason to doubt that an account of incomplete senses requires an attribution of reference, we need look no further than the history of metaphysical inquiry. In the traditional debate about the basic ontological categories, many disagree over the fundamental nature of objects, but few deny that these entities exist or their (rough) character.²⁵ For instance, while there are deep metaphysical questions about the fundamental nature of ordinary objects like tables and chairs, a metaphysician would be hard-pressed to convince a competent speaker of the language that chairs and tables do not exist, or to completely overturn her basic conception of these objects. This conception might include (among other things) assumptions about what it takes for an ordinary object to persist through change, or for two objects to be tokens of the same type, and so on. A subject might, after some prodding, abandon one or more of her assumptions; but it is a strange person indeed who gives up all of them.²⁶

In contrast, real disagreement arises about the existence, nature, and (most importantly) manifest character of those properties instantiated by objects. Are these properties repeatable? Can they exist uninstantiated? Should they be individuated by the causal powers they confer upon objects

²⁵ I have in mind so-called ‘bundle theorists’ who treat objects as appropriate collections of properties (e.g. Williams 1953). More recently, Dasgupta (2009) defends a no-object view (‘generalism’) that constructs facts (rather than objects) from fundamental properties.

²⁶ Cf. Ayers (1991).

that have them? Whereas parallel questions about objects find initial answers in our basic conceptions of objects, these questions about properties do not obviously come with ‘intuitive’ answers. Note that this asymmetry between objects and properties lies at the level of what we might call our ‘common sense scheme’; mature metaphysical theories of each ontological category might prove equally problematic.

We can explain the asymmetry if reference to properties and reference to objects play radically different roles in accounts of the senses of complete and incomplete expressions, respectively. If grasp of the sense of a singular term consists in grasp of a criterion of identity for the object it picks out, the existence and (rough) character of objects will be manifest to a competent speaker of the language. How rich a conception of these objects does a speaker acquire from her grasp of the relevant criteria of identity? It will depend upon how much information our theory of understanding packs into these criteria of identity.

Given what has been said about objects, if grasp of an incomplete expression required grasp of a criterion of identity for properties—a requirement that would make it parallel to Frege’s ‘home case’ of singular terms—we should expect more consensus as to the rough character of these entities. At minimum we should expect to have a feel for what it would take for two properties to be identical (or something near enough). But as we’ve seen, we lack this sort of grasp of the existence or character of properties—unlike in the object case.

These remarks about the history of metaphysical inquiry suggest that we ought to doubt accounts on which grasp of an incomplete sense requires grasp of anything akin to criteria of identity for properties. Yet without this requirement in place, it’s hard to see how a right account of incomplete senses could require an attribution of reference.

1.3.2 An Argument from Fregean Austerity

My second argument operates wholly within the Fregean framework I've sketched over the past few sections. Suppose Frege wanted to be austere. What resources would he need—at minimum—to get his account off the ground? He needs objects to serve as the bearers of names, and truth-values to serve as the referents of sentences. Would he also require functions to serve as the referents of incomplete expressions?

As we've seen, incomplete expressions must have a semantic value. Functions play this role. They capture the pairings of objects and truth-values that incomplete expressions contribute to the determination of the truth or falsity of sentences within which they occur. The assignment of functions as semantic values for incomplete expressions affords Frege a kind of austerity at the level of reference. From a formal point of view, functions are just pairings of collections of objects with truth-values. For example, the function Frege would assign as semantic value for ‘square’ constitutes two pairings of objects: some paired with the True (the square objects), the rest paired with the False (the non-square objects). Functions on this picture are not on an ontological par with objects and truth-values; they are instead a kind of derivative semantic value—a reification of the pairings of objects with truth-values.²⁷

²⁷ These remarks about the possibility of an austere Fregean account at the level of reference lead naturally to an oft-neglected exegetical question about the status of functions within Frege's framework. Dummett assumes—and other commentators have since followed him on this point—that two functions are distinct only if they map different objects to the True (that is, only if they have different ‘extensions’). Some passages in Frege seem to support this assumption (e.g. Frege 1884, §68n); other remarks seem to push in the opposite direction (e.g. Frege FC, p. 26). Of course, Frege *does* distinguish functions from their extensions: the former are incomplete, while the latter are complete (cf. Frege FC, p. 27). Yet this distinction does not alone show that two functions can share the same extension. To establish this stronger result Frege would need grounds upon which to individuate functions in a more fine-grained fashion than extensions. Frege's (1984, §90) emphasis on gap-free proof may provide a starting point. If the substitution of co-extensional incomplete expressions can turn a proof from gap-free to gappy, we may have a relevant difference between functions that cannot be explained by a difference in their extensions. (See Rumfitt 2000 for discussion and development of Frege's notion of gap-free proof.) Stepping back a bit, the fact that Frege's distinction between sense and reference leaves open the possibility of distinct functions that share an extension provides Dummett's challenge even more bite. For if the name/bearer paradigm applied to the relationship between predicates and functions, we would expect a clear answer as to whether functions are individuated by their extensions.

If Frege's account of incomplete senses is compatible with this austere account at the level of reference, the account does not require the attribution of reference to incomplete expressions. For if the account required such an attribution, then functions would have to be more than pairings of objects with truth-values. Recall that Frege admits a more substantive notion of object because of the role objects play in a right account of the senses of singular terms. (I spelled out this connection when introducing the second aspect of Frege's notion of reference.) The purpose of singular terms is to identify a portion of the world for the rest of the sentence to be about. As a result, grasp of the sense of a singular term requires grasp of a criterion of identity for its bearer.

Incomplete expressions cannot play the same role as singular terms within Frege's system. As we've seen, Frege recognised that he cannot solve the UNITY OF UNDERSTANDING PROBLEM if these expressions merely serve to identify an element of the world. To solve the problem he had to recognise the classificatory role of incomplete expressions. Grasp of the sense of an incomplete expression must provide a subject with classificatory knowledge—knowledge of what it would be for the incomplete expression to be true or false of a given object.

This account of incomplete senses—an account driven by Frege's need to solve the UNITY OF UNDERSTANDING PROBLEM—fits with an austere Fregean account at the level of reference. The requisite classificatory knowledge suffices as a basis upon which to pair objects and truth-values. In the case of ‘square’, for example, we know that an object pairs with the True just if ‘is square’ is true of it (and the False otherwise). A thinker need not first single out something in the world—a function—which would in turn determine pairings of objects and truth-values. However this fit between a Fregean account of incomplete senses and the austere story at the level of reference strongly suggests that this account of incomplete senses does not require the attribution of reference to incomplete expressions.

1.3.3 An Argument from Classification

A third argument for the conclusion that a right account of incomplete senses doesn't require the attribution of reference to their associated incomplete expressions concerns the nature of classification. Step back from Frege's framework for a moment. Ask yourself: does someone who understands an incomplete expression exercise anything like the capacity exercised by someone who understands a singular term?

We use singular terms to pick out an aspect of the world for the rest of the sentence to say something about. Incomplete expressions, by contrast, we use to group like with like (e.g. red objects with other red things) and to separate like from unlike (e.g. red objects from green ones). Understanding an incomplete expression must enable a subject to perform this kind of classificatory task. A subject can understand 'tall' or 'red' without *correctly* sorting things as tall or red; what she must possess is sensitivity to circumstances in which it would be appropriate to classify something as tall or red.²⁸ Yet this kind of sensitivity—which manifests classificatory knowledge—does not seem like a mode of identification. In particular, a subject with this sensitivity seems to be in the business of grouping objects together, not isolating or identifying whatever property or feature these objects might have in common. A Fregean account of incomplete senses—an account driven by the UNITY OF UNDERSTANDING PROBLEM—meshes with these independent claims about our understanding of singular terms and incomplete expressions. Yet nothing in the Fregean account or these independent

²⁸ Of course, this sensitivity must be compatible with the various complications that attend the application of natural language expressions to the world. In particular, it cannot run afoul of the vagueness of natural language. As Raffman (2014) persuasively argues, competent applications of vague expressions involve an ineliminable degree of permissible arbitrariness. For example, when told to apply the predicate 'red' to a succession of coloured patches whose colours progress from clear red to clear orange, competent speakers will vary (both inter-personally and intra-personally) with respect to where they stop applying the predicate (or start applying it, if they begin at the orange end of the series). Raffman argues that this variation is faultless: no speaker can count her stopping place as any more (or less) correct than the others. A full account of the classificatory knowledge characteristic of understanding of an incomplete expression must accommodate this feature of vague language (along with many other features of natural language use). Yet what I must say about classifying knowledge remains schematic, and so quite neutral with respect to these further complications.

claims forces us to treat grasp of an incomplete sense as more than a capacity to classify objects in a particular fashion.

These observations about what it takes to understand an incomplete expression do not *prove* that an attribution of reference to incomplete expressions need play no role in a right account of incomplete senses. But this claim about incomplete senses nevertheless becomes quite plausible in light of the arguments developed in these last three subsections. My three arguments have been driven by the thought that incomplete expressions (and thus incomplete senses) play a very different role in our cognitive lives than singular terms (and their senses). In Frege's system this asymmetry has a principled foundation. In order to solve the UNITY OF UNDERSTANDING PROBLEM he must explain grasp of incomplete senses in terms of both grasp of Thoughts and grasp of the senses of singular terms. This explanation seems to force upon him an account of incomplete senses in which our grasp of these senses amounts to mere classificatory knowledge. And what has become clear over the last three subsections is that classificatory knowledge—at least as we've understood it so far—has very little in common with knowledge of a criterion of identity. Yet if we cannot sustain an analogy between the two types of knowledge, we lack grounds upon which to extend the name/bearer paradigm to incomplete expressions.

2 The Generality of the Challenge

I've shown that even if we do not follow Dummett in his verificationism, a version of his challenge still arises within the Fregean framework. But the interest of Dummett's challenge reaches beyond Frege. In this section I argue that a suitably tweaked version of Dummett's challenge applies to a wide range of contemporary accounts of our thought and talk.

Frege's framework set the agenda for contemporary work on mental and linguistic representation. Yet nobody accepts every Fregean commitment. My primary aim in what follows is

therefore to show how much of a departure from Frege would be required to resist Dummett's challenge. Many are caught by the challenge simply in virtue of their doctrinal proximity to Frege.

The section has two parts. In §2.1 I show that the second strand Dummett identifies in Frege's notion of reference—reference modelled on the name/bearer prototype—remains in play even if we adopt a non-Fregean formal semantic theory; the challenge simply re-emerges at the level of thought. In §2.2 I bring out the depth of the UNITY OF UNDERSTANDING PROBLEM.

Together these results exhibit the generality of Dummett's original challenge. It transcends its verificationist origins as a challenge internal to Dummett's Fregean system, and instead applies to a whole range of neo- and non-Fregean attributions of reference to predicative thought and talk.

2.1 The Challenge Generalised I—Mental Representation

2.1.1 On the Task of a Formal Semantic Theory

Dummett was right to recognise two strands in Frege's notion of reference—reference as semantic value and reference modelled on the prototypical relation between a name and its bearer. However the resulting Fregean picture may seem baroque in the face of contemporary work within semantics and philosophy of language.

The history of semantics in the 20th Century could be viewed as a succession of refinements of (and ultimately departures from) Frege's initial semantic apparatus. Semanticists retain Frege's view that the semantic value of a complex sentence should be a function of the semantic values of its constituent expressions. Yet they frequently abandon the Fregean account of these semantic values: the semantic value of a singular term need not be its bearer (or indeed an object at all); the semantic value of a sentence need not be its truth value; and so on.

These departures from Frege have partly resulted from a shift in the dominant conception of the task of a formal semantic theory. I'm going to briefly sketch two somewhat extreme (and necessarily broad-brush) conceptions of this task—one broadly Fregean, the other broadly Lewisian.

I do not take a stand on whether a view in either neighbourhood is correct. My aims in this subsection are more modest. I first highlight how Dummett's challenge remains standing regardless of whether formal semantics makes room for a notion of reference modelled on Frege's name/bearer prototype. I then argue that we should not use the notion of semantic value to replace Frege's more demanding notion of reference. These results are the first step towards generalising Dummett's challenge beyond its initial Fregean setting.

Dummett read Frege as forging a tight link between our formal semantic theory and our theory of understanding (a.k.a. Frege's theory of sense). Suppose you wish to provide a theory that captures the rational position of someone who speaks the language. What you want to do is match up expressions with what they stand for in a way which recognises differences between expressions which stand for the same thing. And if you write down a theory that does that, you will have a theory whose clauses state the referents of expressions and show their senses. And that is going to be a Fregean semantic theory.²⁹

Frege puzzles, for instance, reveal the difference in rational position between someone who understands 'Cicero is Cicero' and someone who understands 'Tully is Cicero'. What Frege recognised was that a right account of the rational difference between accepting 'Cicero is Cicero' and accepting 'Tully is Cicero' must exploit facts that concern both what 'Cicero' and 'Tully' stand for—namely Cicero—and how these ways of picking out the man differ (in particular, someone who understands 'Cicero' and 'Tully' grasps different modes of presentation of Cicero).

²⁹ For clear statements of this conception of a semantic theory, see Dummett (1973/1981, p. 227; 1991b, pp. 237–8) and Evans (1982, pp. 25–7). Here is Dummett (1973/1981): 'Indeed, even when Frege is purporting to give the sense of a word or symbol, what he actually *states* is what the reference is: and, for anyone who has not clearly grasped the relation between sense and reference, this fact makes his hold on the notion of sense precarious. The sense of an expression is the mode of presentation of the referent: in saying what the reference is, we have to choose a particular way of saying this... In a case in which we are concerned to convey, or stipulate, the sense of an expression, we shall choose that means of stating what the referent is which displays the sense: we might here borrow a famous pair of terms from the *Tractatus*, and say that, for Frege, we *say* what the referent of a word is, and thereby *show* what its sense is.'

Dummett, following the lead of Frege, held that a formal semantic theory should assign semantic values to expressions in a way that helps capture these kinds of facts about speaker understanding—facts that ultimately limn the rational life of those who use and understand the language.³⁰ It was this conception of the task of a formal semantic theory that was partly responsible for securing a role within semantics for a conception of reference modelled on Frege's name/bearer prototype.

A broadly Lewisian strand of contemporary work in semantics departs from this Fregean conception of the task of a formal semantic theory. Instead of treating semantic theories as importantly answerable to our theory of understanding—to facts about the rational position of those that understand the language—followers of the Lewisian approach see the task of a formal semantic theory as almost exclusively one of capturing the internal compositional and inferential structure of natural language.³¹ These theorists will deny that a semantic theory must capture facts about cognitive significance (the facts exploited by Frege Puzzles). If such facts are to be accommodated at all, this accommodation will not find direct reflection in the semantics of a particular language. For example, one might explain a difference in cognitive significance between ‘Hesperus is bright’ and ‘Phosphorus is bright’ by appeal to a metasemantic divide between ‘Hesperus’ and ‘Phosphorus’ rather than by appeal to any distinction at the level of semantics.

Another way of putting the contrast between these two conceptions of semantics is in terms of the constraints each side respects when constructing a semantic theory.

A semantic theory must satisfy certain constraints when assigning semantic values to atomic expressions, and specifying rules for how these semantic values combine to determine the semantic

³⁰ See also Rattan (2010, §2.2) for a perspicacious statement of this connection between sense, rationality, and reference.

³¹ Cf. Lewis (1970) and Heim and Kratzer (1998). Of course, those I've been calling ‘Lewisians’ don't ignore the ‘Frege-Dummett’ demand to explain facts about rationality and understanding. Many follow Dummett's own lead in distinguishing a notion of content relevant for rational psychology (‘assertoric content’) from a notion of content relevant for a compositional formal semantics (‘ingredient sense’). For discussion of this oft-neglected distinction, see Dummett (1973/1981); Dummett (1991a); Lewis (1980a); Stanley (1997); Yalcin (2007); Rabern (2012).

values of complex expressions. What I've been calling 'Lewisian' semanticists recognise only purely formal constraints (to generate semantic values for complex expressions that receive the correct truth-values relative to contexts; to assign uniform semantic values within a syntactic category; to capture certain entailments between sentences; and so on). In contrast, members of the Fregean camp insist on further demanding rationality-involving constraints on the construction of a semantic theory (for example, somebody in this camp will say that differences in the circumstances under which competent speakers of the language would rationally accept as true the sentences 'Cicero is wise' and 'Tully is wise' entail differences in right treatments of the expressions 'Cicero' and 'Tully' in a semantic theory).

The dispute between these broadly Fregean and Lewisian conceptions of semantics therefore partly concerns whether the notion of reference—as modelled on the name/bearer prototype—ought to play a role in our formal semantic theory. The formal semantic study of natural language is driven by methodological principles that potentially run afoul of our intuitive conception of aboutness. For example, semanticists generally adhere to a uniformity constraint that parallels principles of parsimony in natural science. The constraint holds that whenever feasible, expressions of the same syntactic type should be assigned semantic values of a single semantic type. Hence a semanticist interested in demonstrative uses of 'that' may also wish her account to also capture the use of 'that' within propositional attitude attributions (e.g. 'Guillam believes that Smiley snored'). But such uniformity principles could very easily lead a Lewisian to abandon reference as a semantically significant relation. For it will not be obvious that the formal constraints governing a 'Lewisian' semantic theory will prevent uniformity principles (or some other formal constraint) from sidelining reference as modelled on the name/bearer prototype. Yet even if a Lewisian permits reference a role in semantics, she can deny the kind of link between a theory of sense and the attribution of reference to a class of expressions which is explicit in 3 (recall that 3 is the claim that an attribution of reference to a class of expressions is justified only if it is required by a right account

of their senses). Many philosophers have become sceptical that there exists a close tie between a language's referential properties and a speaker's epistemic position. A speaker need not—or so the line of thought goes—have any substantial epistemic contact with the referent of an expression she understands. Some of this pressure towards a sharp break between the epistemic and the semantic comes from the public dimension of language.³² Speakers of a natural language are able, merely through deference to the use made by a community of speakers, to extend their ability to express information about the world. And it seems like a small jump from such linguistic 'outsourcing' to a view on which our semantic reach far outstrips our epistemic limitations.

However notice that the Lewisians need not deny the importance of a notion of reference modelled on the name/bearer prototype; they simply believe semantics does not need it. Someone who wishes to properly resist Dummett's challenge must make the stronger claim that reference *never* plays an important explanatory role with respect to speaker understanding. For as we saw in §1.1.2, it is this link between reference and understanding that underwrites the name/bearer prototype, not any role for reference within formal semantics (for instance within a right account of the truth conditions of sentences that include first or second order quantifiers). Later in §2.1.3 I argue that when we opt for the Lewisian approach to semantics, the link between reference and understanding finds an alternative home at the level of thought (a variety of mental representation).

2.1.2 An Informative Objection

At this point a determined opponent may object: 'why tie a theory of sense to the name/bearer prototype, when we might instead hold that the sense of an expression is a mode of presentation of its semantic value?' In other words, why not combine our formal semantics and theory of understanding together, and dismiss as baroque Frege's name/bearer prototype? For this objection

³² Cf. Kaplan (1989, pp. 603–7). Kaplan's framework provides a nice example of the Lewisian approach. For on that framework, questions about cognitive significance—the province of Frege's notion of sense—largely belong to 'metasemantics' rather than to semantics proper (Cf. Kaplan 1989, 573–4).

to constitute a genuine challenge, these opponents must reject the Fregean identification of semantic value with the bearer of a name, as well as parallel identifications for other referring expressions.

This objection obviously isn't available to those sympathetic to what I've called a 'Lewisian' conception of the task of a formal semantic theory. For we've seen that Lewisian semanticists take their task to be largely independent of concerns about the nature and structure of understanding. But this independence cannot be sustained if we chain our theory of sense, and hence our theory of understanding, to our formal semantics. Semantic value assignments would become sensitive to precisely the kinds of facts that a Lewisian denies a formal semantic theory must capture (namely facts about the rational positions of those who understand the language).

I think the objection also fails on non-Lewisian grounds. However it remains important for what its failure reveals. As we'll see, the objection's failure reveals the relative independence of a Fregean theory of understanding from a purely formal semantics. I shall argue that to chain the theory of sense to a pure theory of semantic value (as our hypothetical objector proposes) generates some bad results for our theory of understanding. In particular, tying our theory of understanding (Frege's theory of sense) to a formal theory of semantic value would violate attractive theses about the structure of understanding.

First, there is an old Fregean intuition that the world imposes substantial constraints on what is required to think about it. One version of this intuition holds that there must be some kind of match between the structure of thought and the structure of that part of the world that would make it true (the associated Thought's truthmaker).³³ For example, this match could derive from the ontological categories to which the referents of the constituents of the thought belong. However purely formal constraints on the construction of a semantic theory permit assignments of semantic

³³ This version of the intuition can be found implicitly or explicitly in the work of many thinkers. For example: Frege (FC; SR), Russell (1913/1984), Wittgenstein (1921), Strawson (1959), Evans (1982), and Dickie (2010b).

values to sub-sentential expressions that do not reflect the structure of the real-world entities that ground the truth of the corresponding sentences (the truthmakers of the sentences). For example, on Generalised Quantifier Theory the semantic value of the proper name ‘Andrew’ is the set of all sets containing Andrew, and the semantic value of the predicate ‘is red’ is the set of all red objects.³⁴ Thus the sentence ‘Andrew is red’ is true iff the set of red objects is a member of the set of all sets containing Andrew (more formally: ‘Andrew is red’ is true iff $\{x \mid x \text{ is red}\} \in \{X \mid \text{Andrew} \in X\}$). Meanwhile, intuitively the ground for the truth of this sentence is the fact that Andrew has the property of being red (or, in Frege’s terms, that Andrew falls under the concept *red*). Given this feature of purely formal semantic value assignments, our old intuition becomes threatened if a Thought is composed of modes of presentation of semantic values. For the structure of our grasp of a thought (one with the Thought as content) may not match the structure of the ground of its truth. So to chain a theory of sense to a theory of semantic value would violate a potentially plausible constraint on the structure of thought.

Second, Fregeans find it plausible that the structure of understanding must be determinate. Yet some assignments of formal semantic value possess an inherently indeterminate structure. For example, recall that on Generalised Quantifier Theory the semantic value of ‘Andrew’ is the set of all sets containing the bearer of ‘Andrew’. This set is in turn formally equivalent to the function denoted by $\lambda P[P(\text{Andrew})]$. As a result, our purely formal semantic theory cannot tell us whether semantic values are functions or sets. Therefore if we chain our theory of sense to a broadly Lewisian theory of formal semantic value, and if the structure of our understanding of sentences is sensitive to the kind or category of the entities presented by the senses of a sentence’s sub-sentential expressions, we must deny understanding a determinate structure.

³⁴ Barwise and Cooper (1981, pp. 164–9).

These last two subsections have explored whether Dummett's challenge remains standing if we abandon a Fregean conception of the task of a formal semantic theory. I've argued that adopting an anti-Fregean 'Lewisian' conception of semantic theorizing merely excludes from semantics proper a notion of reference modelled on the name/bearer prototype; it does not eliminate the need for such a rich notion of reference altogether. (Indeed in this subsection I've highlighted the cost of eliminating the rich notion of reference.) In the next subsection I show how Dummett's challenge emerges at the level of thought even if we adopt a Lewisian conception of semantics.

2.1.3 The Challenge for Mental Representation

Now that I've argued that the central component of Dummett's challenge remains even if we abandon a Fregean conception of the task of a formal semantic theory, it will be useful to have a reformulation of the challenge as a challenge about thought (a variety of mental representation). The challenge arose originally because of Frege's apparent commitment to 1–4:

1. All (non-degenerate) significant expressions possess both sense and reference.
2. The senses of incomplete expressions are themselves incomplete (or 'unsaturated').
3. An attribution of reference to a class of expressions is justified only if it is required by a right account of their senses.
4. An account of the incomplete sense of an incomplete expression does not require the attribution of reference to the expression.

Dummett raised the challenge in the way he did because he followed Frege in thinking that language and Thought go hand in hand. He believed that a sentence expresses a Thought; that to understand a sentence is to grasp a Thought; and that the way to talk about linguistic expressions is in terms of their contribution to Thoughts. But I've argued over the last few pages that the challenge remains even if we sever this tight Fregean link between language and Thought. As we'll see, the challenge still arises at the level of thought.

Let's fix some ideas in order to triangulate Dummett's challenge for right accounts of thought. I assume that to think a thought one might express with a sentence of the form ' α is Φ ' involves the exercise of at least two distinct abilities: an ability to think about the referent of α and an ability to think that something is Φ . The latter ability is the mental analogue of the classificatory capacity possessed by someone who understands an incomplete expression. If we individuate abilities in terms of the conditions of their exercise, we can define a 'concept' as a reification of these conditions (not to be confused with Frege's notion of a concept as a function from objects to truth values).³⁵ Hence to think the thought one might express with a sentence of the form ' α is Φ ' is to token a conceptual representation composed of at least those concepts whose possession amounts to the exercise of abilities to think about the referent of α and to think that something is Φ . We represent this conceptual representation—a thought—with ' $\langle\alpha \text{ is } \Phi\rangle$ ', using ' $\langle\alpha\rangle$ ' and ' $\langle\Phi\rangle$ ' to denote the two constituent concepts.

Concepts may be made as fine-grained as Fregean senses, since a capacity to think about an object, for example, can be exercised in many different ways. Each way of thinking about an object will require a thinker to satisfy distinct conditions, and for each such set of conditions we can define a concept.

The shift to thought requires a reformulation of Dummett's challenge. In our original formulation, expressions occupied a central position in virtue of their role as bearers of sense and reference. But now we need not distinguish the role of sense and of bearing reference: concepts can do both jobs. So the core of Dummett's challenge becomes the following:

³⁵ For definitions of concepts with this general structure, compare Evans (1982), Peacocke (1992; 2008), Stanley (2011), and Dickie (2012). This isn't a neutral definition of 'concept', but what I intend to say about thought does not ultimately depend upon how exactly we make sense of concepts as components of thoughts. As we'll see, Dummett's challenge arises for any account of thought that treats concepts as the bearers of reference and also affords them a role in the determination of cognitive significance. However I do mean to distinguish my notion of 'concept' from what psychologists usually call 'concepts', namely mental categories used by thinkers to group objects, events, and the like. While there are difficult questions about the relationship between concepts as components of thoughts and concepts as mental categories, only confusion can follow an identification of the two notions.

1. All (non-degenerate) concepts possess both determinate possession conditions (a determinate sense) and reference.
2. Predicative concepts are incomplete.
3. The attribution of reference to a class of concepts is justified only if this attribution is required by a right account of what it is to possess these concepts.
4. A right account of the possession conditions of incomplete concepts does not require the attribution of reference to these concepts.

A suitably tweaked version of the UNITY OF UNDERSTANDING PROBLEM motivates the introduction of predicative incomplete concepts (and thus 2), and the arguments that motivated the original version of 4 apply equally to this new version. Yet as before, 2–4 entail that an attribution of reference to predicative concepts would be unjustified. This is the version of Dummett's challenge that generalises to a whole range of contemporary views which attribute reference to predicative concepts (those views which include commitment to something like 1).³⁶

2.2 The Challenge Generalised II—The Depth of the UNITY OF UNDERSTANDING PROBLEM

Once we accept a link between our theory of reference and our theory of understanding (whether at the level of language or the level of thought), the other obvious way to cut Dummett's challenge off at the knees would be to reject a suitably updated version of the problem that both motivated and constrained Frege's account of incomplete senses:

UNITY OF UNDERSTANDING PROBLEM: The thought $\langle a \text{ is } F \rangle$ is composed of ways of thinking of **a** and **F**. A thinker who grasps $\langle a \text{ is } F \rangle$ counts as thinking that **a** is **F**.

³⁶ To name just a few vulnerable accounts: Evans (1982); Peacocke (1992); Chalmers (2006b; 2012); Wedgwood (2007); Camp (2009).

The alternative is that she counts as thinking of a mere list of entities (i.e. the list consisting of **a** and **F**). In virtue of what does she count as thinking that **a** is **F**?

This problem does a lot of work within Dummett's challenge. It motivates both the extension of the complete/incomplete distinction to the level of concepts, and the claim that to grasp an incomplete predicative concept **<F>** just is to know what it would be for **<F>** to be true or false of a given object (a claim largely responsible for 4: the thesis that a right account of the possession conditions of incomplete concepts does not require the attribution of reference to these concepts).

In this subsection I explore the depth of the UNITY OF UNDERSTANDING PROBLEM. I first argue that it is distinct from (and arguably more fundamental than) traditional metaphysical problems concerning so-called 'propositional unity'. I then show that my problem afflicts even those who avoid traditional metaphysical unity problems by denying that propositions have internal structure.

2.2.1 Russell Understood

Propositions in this context are abstract entities that (many assume) serve as the semantic content of sentences and intentional states like belief. The sentence 'Lyra loves Will' therefore expresses the same proposition that Bill believes when he believes that Lyra loves Will. In the case of sentences, propositions are partly determined by the semantic values of sub-sentential expressions. Consider neo-Russellian views on which the semantic value of a name is the object it stands for, and the semantic value of a relational expression is a relation. On this view the proposition expressed by 'Lyra loves Will' is composed of Lyra, Will, and the relation of loving.

Russell set the agenda for discussion of what has become known as the 'problem of propositional unity'. As he recognised, the so-called 'problem' actually encompasses a list of problems that plague almost any theorist who wishes to introduce propositions. Perhaps most obvious is the problem of explaining how constituents of a proposition combine to constitute a

whole proposition rather than a mere collection of constituents. For a neo-Russellian this becomes the problem of explaining how Lyra, Will, and the relation of loving (to continue with the same example) constitute a whole proposition rather than a mere collection of entities. By contrast, a neo-Fregean for whom propositions are composed of modes of presentation must explain how the relevant modes of presentation (in our example, modes of presentation of Lyra, Will, and the relation of loving) together combine to constitute a proposition.

A closely related issue concerns internal propositional structure: if ‘Lyra loves Will’ expresses a proposition, it is surely different from the proposition expressed by ‘Will loves Lyra’. Yet if propositions have constituents, surely these two share their constituents (regardless of whether these constituents are the neo-Russellian’s objects and properties, or the neo-Fregean’s modes of presentation). But what else could account for the difference between the propositions expressed by the two sentences, if not their constituents?³⁷

A third problem has more recent provenance. It asks how an entity composed of propositional constituents (for example, Lyra, Will, and the relation of loving) could *represent that* Lyra loves Will. More generally, it is the problem of explaining how a proposition composed of sub-propositional constituents could possess truth conditions.³⁸

The contemporary debate over the various problems of propositional unity largely ignores my UNITY OF UNDERSTANDING PROBLEM. Yet Russell saw the problem even more clearly than Frege did. I take Russell to have the problem in mind when he observes that

In order to understand “A and B are similar”, we must know what is supposed to be done with A and B and similarity, i.e. what it is for two terms to have a relation; that

³⁷ For the first puzzle, see Russell (1913/1984, p. 116); for the second, see Russell (1913/1984, pp. 86, 111).

³⁸ Eklund (ms) attributes the original problem to Sainsbury (1996). King (2007) and Soames (2010) both resolve the problem by denying one of its presuppositions: the claim that propositions possess intrinsic truth conditions (namely truth conditions that do not derive from more basic linguistic or mental facts).

is, we must understand the form of the complex which must exist if the proposition is true.³⁹

Whereas the ‘traditional’ unity problems concern the metaphysics of propositions—entities introduced to serve a variety of explanatory roles—my UNITY OF UNDERSTANDING PROBLEM arises even before we introduce propositions (recall that in order to solve this problem Frege introduced the complete/incomplete distinction at the level of sense). For regardless of whether anything satisfies the explanatory roles propositions are standardly taken to occupy, there remains a question about what it takes for a subject to understand what is said and what is thought.

2.2.2 Unity Without Structure

Someone who denies propositions an internal structure arguably avoids the problem of explaining the compositional unity of propositions—how their elements ‘compose’ or ‘hang together’ to form a single entity. Stalnaker (1984), for example, takes propositions to be sets of possible worlds. On this view the proposition expressed by ‘Lyra loves Will’ is just the set of possible worlds in which Lyra loves Will. Sets have no parts, only members, and so the traditional metaphysical problem seemingly does not arise.

Can Stalnaker (and others like him) take a similarly dismissive attitude to my Fregean UNITY OF UNDERSTANDING PROBLEM? If they can, Dummett’s challenge is easily resisted. But I contend that a version of the UNITY OF UNDERSTANDING PROBLEM arises even within a Stalnakerian framework, and so there remains a hard question about whether he (and others like him) can resist its force.

The UNITY OF UNDERSTANDING PROBLEM targets Thoughts. These occupy an explanatory role distinct from that occupied by the propositions targeted by the traditional metaphysical unity problems. Stalnaker’s treatment of propositions follows from his account of what a thinker is able to

³⁹ Russell (1913/1984, p. 116).

do when she understands a sentence.⁴⁰ A thinker understands a sentence S iff she has the capacity to divide the space of possibilities in the right way (in the case of S, a capacity to discriminate between worlds in which it is true and those in which it is false). Propositions are therefore just ways of carving up the space of possibilities.

Stalnaker thus models propositions only on the output of our capacity to grasp them (our capacity to understand sentences which express these propositions). But such an account does nothing to illuminate how understanding works, what we might call the ‘mechanism of understanding’. It is this lacuna that Fregean propositions (*Thoughts*) occupy. Fregean propositions are composed of ways of thinking about objects and properties. The fact that a thinker engages in these ways of thinking, and therefore possesses knowledge of what it would be for the thought to be true, explains her capacity to discriminate between possibilities in the way Stalnaker demands. Fregean propositions thus seek to capture the internal structure of understanding, while Stalnakerian propositions seek to capture the structure of the output of understanding. These two projects are clearly distinct, and so Stalnaker’s way out of the traditional metaphysical unity problem cannot straightforwardly extend to the **UNITY OF UNDERSTANDING PROBLEM**.

Of course Stalnaker could deny that the internal structure of understanding is as rich or determinate as the Fregean framework suggests.⁴¹ But such a denial might not save him from the **UNITY OF UNDERSTANDING PROBLEM**. Take the sentence ‘Alexander Pope poisoned Edmund Curl’. On Stalnaker’s framework, a thinker understands this sentence iff she can discriminate the worlds in which Alexander Pope poisoned Edmund Curl from those in which he did not, and knows in which set the sentence is true. But then it makes sense to ask: in virtue of what does the thinker discriminate the worlds in which Alexander Pope poisoned Edmund Curl from the worlds in

⁴⁰ Stalnaker (1984, pp. 4–5).

⁴¹ As it turns out, Stalnaker (1981, p. 134) seems to endorse the attribution of a rich internal structure to understanding.

which Alexander Pope and Edmund Curl both exist, and in which poisonings occur, but in which neither man poisons the other? It is precisely this kind of question to which the UNITY OF UNDERSTANDING PROBLEM demands an answer.

3 A Landscape of Responses

With Dummett's challenge now fully on the table, I want to quickly canvas a number of potential responses. The survey will not be exhaustive. Which response we ought to pursue very much remains an open question at this stage. Here are the central components of the challenge once again at the level of thought:

1. All (non-degenerate) concepts—where these are components of thoughts rather than Fregean functions—possess both determinate possession conditions (a determinate sense) and reference.
2. There are incomplete concepts.
3. The attribution of reference to a class of concepts is justified only if this attribution is required by a right account of what it is to possess these concepts.
4. A right account of the possession conditions of incomplete concepts does not require the attribution of reference to these concepts.

Our basic options are clear: avoiding Dummett's challenge requires rejecting one (or more) of 1–4. What might these responses look like?

Reject 1: This is the nuclear option. As I've developed it, 1 more or less falls out of a highly attractive and intuitive conception of thought on which the formation of beliefs—and thoughts more generally—requires the exercise of capacities to represent objects and their properties. 1 requires little of these representational capacities: only that they must combine to enable subjects to think, of an object, that it is a certain way. Hence 1 carries no obvious commitment to controversial

claims about thought more generally, or about representational capacities in particular. Someone tempted to reject 1 therefore faces the daunting task of dislodging an otherwise anodyne conception of how the mind represents things outside itself. In particular, to reject 1 a theorist must do more than simply hive off a particular class of exceptional expressions (or concepts) and insist (as one might in the case of empty names) that members of this class fail to refer yet nevertheless contribute to the truth or falsity of sentences in which they occur.⁴² Opponents of 1 must instead defend the stronger claim that there is this whole class of expressions or concepts—the predicative or ‘incomplete’ ones—which don’t do their ‘hooking up with the world’ work in virtue of standing for some kind of entity.⁴³

Reject 2: We’re motivated to accept 2, as well as its analogue for language, largely because these claims seem to help solve their respective versions of the UNITY OF UNDERSTANDING PROBLEM. Given this background, there are two grounds upon which one might legitimately reject 2: either the claim fails to help secure a viable solution to the UNITY OF UNDERSTANDING PROBLEM, or there exists a better solution to the problem that does not require the problematic sort of distinction between complete and incomplete concepts (or more generally between complete and incomplete senses). For example, in the second half of *Individuals* P.F. Strawson draws a complete/incomplete distinction at the level of language in terms a difference between the presuppositions carried by singular terms and those carried by predicative expressions (namely: only our uses of singular terms ‘carry a weight of fact’ by committing us to something about the world).⁴⁴ However I will not pursue this alternative here. My aim has been to explore Dummett’s challenge as it arises within a Fregean

⁴² I have in mind those like Braun (1993) who argue that sentences containing empty names are false (rather than truth-valueless) in virtue of expressing defective or ‘gappy’ propositions. These theorists insist that the referent of a name is its bearer, but allow that the absence of a referent also constitutes a contribution to the truth or falsity of sentences within which the name occurs.

⁴³ Notice that even adopting nominalism about properties would not yet secure this stronger claim. Nominalists can allow that incomplete expressions or concepts stand for extra-linguistic entities; they merely believe that positing these entities does not commit us to anything over and above an ontology of objects.

⁴⁴ Strawson (1959, pp. 180–7).

framework where the understanding of the complete/incomplete distinction is as I've described it. Strawson's version of the distinction, by contrast, represents a quite different line of exploration.

Reject 3: I've shown that even those who banish Fregean concerns about rationality from their formal semantic theory—and with them also the analogue of 3 for language—face 3 at the level of thought. For our commitment to 3 seems to fall out of the very minimal conception of thought that underwrites 1, a conception that individuates concepts (like Frege individuates senses) partly in terms of their fine-grained roles within our rational cognitive lives.

However 3 does face potential opposition from another quarter. The role of 3 within Dummett's challenge is to tie our theory of sense to our theory of reference. It must do so in such a way that our account of the possession conditions of incomplete concepts does not require the attribution of reference to these concepts (that is, 3 must make 4 plausible). Yet 3 could remain true, and forge an important connection between our accounts of sense and reference, without thereby lending support to 4. In particular, someone armed with an alternative account of 'complete' concepts—concepts of objects—might dislodge Dummett's appeal to a subject's grasp of criteria of identity. Yet this manoeuvre will answer Dummett's challenge only if the new account of 'complete' concepts undermines 4 without in the process running afoul of the incompleteness of incomplete concepts.

Reject 4: The plausibility of 4 has two sources. First, it's quite unclear what sort of (otherwise workable) account of incomplete concepts could capture the necessary parallel to Frege's name/bearer prototype. Suppose we follow Dummett in treating thought about objects as bound up with grasp of criteria of identity. What sense can we make of the claim that possession of an incomplete concept involves grasp of a criterion of identity for observable properties or (even more bizarrely) Fregean functions? It was this question that drove the first of my arguments for the original version of 4 back in §1.3 (the argument from the history of metaphysics).

The need to solve the UNITY OF UNDERSTANDING PROBLEM imposes constraints on a right account of incomplete concepts. These constraints also contribute to the plausibility of 4. We've seen that solving the problem usually demands little of the 'referential' dimension of incomplete concepts (or senses); instead a subject must manifest an appropriate sort of classificatory knowledge (knowledge that seems a far cry from the identifying knowledge speakers must have in order to possess complete concepts). As a result, if the central explanatory demands on incomplete concepts (and senses) come from the UNITY OF UNDERSTANDING PROBLEM, we have few resources with which to resist 4. For how can we insist that a right account of incomplete concepts requires an attribution of reference to them, when this attribution seems surplus to requirements with respect to the unity of understanding? It was this thought that underwrote my other two arguments for the original version of 4 back in §1.3 (the arguments from Fregean austerity and from classification).

Given this background, opponents of 4 must accomplish two tasks: (1) explain how an attribution of reference to incomplete concepts can figure in an account of their possession conditions; (2) show either (a) that such an account actually follows from the explanatory demands imposed by the UNITY OF UNDERSTANDING PROBLEM⁴⁵ or (b) that such an account follows from some other set of explanatory demands.

4 Conclusion

My primary goal has been to show that even those who have taken major steps away from Frege's framework cannot ignore Dummett's challenge. The challenge remains even if one abandons a

⁴⁵ This reply parallels a well-known style of objection to deflationism about truth. Deflationists often characterise their view as follows (cf. Field 1994): the instances of a chosen equivalence schema for truth (e.g. 'p' is true iff p) exhaust what can be significantly said about truth. Crispin Wright (1994) objects that a concept of truth satisfies the equivalence schema only if truth serves as a norm for assertion (and thus the equivalence schema forces us to say something quite substantive about truth). But if truth must serve as the basis of a substantive norm, Wright argues, deflationism cannot be sustained. Wright's objection remains controversial. For example, Rumfitt (1995) argues that a deflationist appeal to a role for denial could allow them to deny a substantial truth norm for assertion.

Fregean conception of formal semantic theory, or insists that mental (or linguistic) content lacks internal structure.

I have not shown—nor have I sought to show—that the challenge has no solution within a broadly Fregean framework. Indeed, I develop my own solution to the challenge as a conclusion to the dissertation. It will exploit resources made available by Chapters 1 and 2.

My exploration of Dummett’s challenge also remains silent about those who pursue more radical departures from Frege’s framework. Some of these departures will open up positions for which Dummett’s challenge fails to arise.⁴⁶ However an investigation of the prospects for such radical non-Fregean alternatives lies outside the ambit of this chapter.

⁴⁶ I’ll sketch one position of this sort. Suppose one follows Fodor (1987) and insists that to think *a* is *F* is just to token an appropriate syntactically structured mental representation in the brain. Add to this ‘language of thought’ thesis a minimal account of understanding: to understand *a* is *F* is just to token a mental representation that represents that *a* is *F*. On the resulting framework my UNITY OF UNDERSTANDING PROBLEM becomes the problem of explaining how a tokened mental representation can count as representing that *a* is *F* (as opposed to merely representing a list of entities). King (2007) solves the analogous problem for sentences by appeal to the (supposedly) rich and determinate syntactic structure present in natural language. For example, King claims that we can use the sentence ‘Buttercup loves Westley’ to represent that Buttercup loves Westley in virtue of the complex fact that (1) Buttercup, the relation of loving, and Westley stand in representational relations to (respectively) ‘Buttercup’, ‘loves’, and ‘Westley’; and (2) these three expressions instantiate a determinate Chomskian syntactic structure (a structure which permits, *inter alia*, a distinction between what ‘Buttercup loves Westley’ and ‘Westley loves Buttercup’ represent). Given that the Fodorian language of thought contains representations with their own syntactic structure, one could try to use an analogue of King’s view to solve the UNITY OF UNDERSTANDING PROBLEM (without positing anything like incomplete concepts). Yet for this extreme position to become available one must first establish some controversial empirical claims about the mind. In particular, one must show not only that thought works as Fodor says it does, but also that the kind of syntactic structure present in an (empirically acceptable) language of thought is rich enough to support an analogue of King’s account. And this is a big ask.

Conclusion: Another Variety of Reference

I shall close by showing how the resources of Chapters 1–2 generate a solution to the fully general version of Dummett's Challenge developed in Chapter 4. The challenge sought to undermine the claim that we think and talk about properties. At the level of thought the challenge arises because of a tension between the following four claims:

1. All (non-degenerate) concepts possess both determinate possession conditions (a determinate sense) and reference.
2. There are incomplete concepts.
3. The attribution of reference to a class of concepts is justified only if this attribution is required by a right account of what it is to possess these concepts.
4. A right account of the possession conditions of incomplete concepts does not require the attribution of reference to these concepts.

Material from Chapters 1–2 allows me to reject 4. I show that something other than the UNITY OF UNDERSTANDING PROBLEM motivates an alternative account of a subset of incomplete concepts. My alternative account carves out a role for the attribution of reference to these concepts that permits me to reject 4. I shall now explain how.

Recall that Chapter 2 developed an argument for PREDICATIVE ATOMISM. The argument begins with a claim about the relationship between our capacity to ascribe unobservable properties (like spin or charge) and our capacity to ascribe causal relations:

1. CAUSALISM: Our capacity to ascribe unobservable properties exploits our capacity to ascribe the causal relations that hold in virtue of these properties (i.e. causal relations that constitute the causal roles of these properties).
2. CAUSAL ATOMISM: Our capacity to ascribe causal relations that hold in virtue of unobservable properties type-depends upon our capacity to ascribe special causal relations.

Special causal relations are those that involve the ordinary objects we encounter in perceptual experience: a boulder *flattens* a hut, a child *kicks* a ball, a cat *slurps* her water, and so on. 1 and 2 together entail 3:

3. Our capacity to ascribe unobservable properties type-depends upon our capacity to ascribe special causal relations. [From 1 and 2]
4. CONJECTURE: Our capacity to ascribe special causal relations type-depends upon our capacity to perceptually demonstrate objects (i.e. to identify objects solely on the basis of perceptual encounters with them).

CONJECTURE is Campbell's (2002) answer to a puzzle Shoemaker (1988) raises for our conception of ordinary objects. Given 3 and 4, 5 follows:

5. Our capacity to ascribe unobservable properties type-depends upon our capacity to perceptually demonstrate objects. [From 3 and 4]

I used 5 as part of an argument for 6:

6. BRIDGE PREMISS: Our capacity to ascribe special causal relations also type-depends upon our capacity to perceptually ascribe observable properties as categorical.

Recall that a subject ‘perceptually ascribes a property as categorical’ iff perception permits her to ascribe the property both (1) without relying only upon knowledge of causal powers it confers, and (2) by treating the property as a ground of an object’s dispositions (e.g. someone who ascribes circularity as categorical can make sense of a circular object being inclined to roll on smooth inclines). In any case, 3 and 6 together entail

7. PREDICATIVE ATOMISM: our capacity to ascribe unobservable properties type-depends upon our capacity to perceptually ascribe observable properties as categorical. [From 3 and 6]

My main contribution to this argument for PREDICATIVE ATOMISM was my defence of the BRIDGE PREMISS. I argued for the BRIDGE PREMISS by defending three nested component claims: (1) our capacity to ascribe special causal relations type-depends upon our capacity to ascribe observable

properties; (2) this ascription must be perceptual; and (3) we must ascribe at least some of the properties as categorical.

Given that we do ascribe unobservable properties, PREDICATIVE ATOMISM entails that we perceptually ascribe at least some observable properties as categorical. Yet it's also the case that the BRIDGE PREMISS carves out an important explanatory role for our capacity to perceptually ascribe observable properties as categorical. For when taken together with the rest of my argument for PREDICATIVE ATOMISM, the BRIDGE PREMISS entails that a right account of our capacity to perceptually ascribe properties as categorical must accommodate its role in permitting us to ascribe special causal relations. And this represents a substantial constraint on a fundamental level of predicative thought.

Given this new constraint, must a right account of our capacity to perceptually ascribe properties as categorical treat this ascription as another variety of reference? Briefly: 'yes'.

In my argument for the BRIDGE PREMISS I made much of the fact that our capacity to ascribe properties as categorical underwrites our ability to isolate (in thought) the changes characteristic of particular special causal relations (see §3.2.1 and §3.2.3 in Chapter 2). These are changes a subject must isolate (in thought) in order to ascribe the corresponding special causal relations.

Yet in order to isolate a change, we must be able to ascribe a higher order property to those properties we exploit when isolating the change. For when we isolate a particular type of change (say a change in an object's shape from round to oval) we distinguish it from other potential changes (such as a change in shape from oval to square). Distinguishing a change in this way requires grouping properties (such as certain shape and size properties) in a way that will individuate the change (for example, an object counts as *crushed* when its new shape is much more flat—and its new size much smaller—than the object's original shape and size). And I argued in Chapter 1 that this

kind of classification requires an ability to bring the items classified together under a common property—in this case a higher order property of properties.

However being able to think that a property has a higher order property requires that a subject be able to isolate the former property as the entity upon whose properties the truth of the thought depends. For just as understanding a sentence of the form ' α is Φ ' requires knowing what it would be for the object α stands for—its referent—to satisfy Φ , thinking $\langle\Phi\text{ is }\Gamma\rangle$ requires knowing what it would be for *the property* $\langle\Phi\rangle$ *stands for* to satisfy $\langle\Gamma\rangle$. (For example, just as understanding ‘Cromwell is asleep’ requires knowing what it would be for Cromwell to satisfy ‘asleep’, thinking $\langle\text{Being asleep is a property of animate objects}\rangle$ requires knowing what it would be for the property of being asleep to satisfy the predicative concept $\langle\text{a property of animate objects}\rangle$.) As a result, in showing that perceptual ascription of a property as categorical must *already* enable a subject to think that the property is Γ —for some appropriate range of higher order properties—I’ve shown that thinking $\langle o \text{ is } \Phi \rangle$ already requires knowledge of what it would be for *the property* $\langle\Phi\rangle$ *stands for* to satisfy $\langle\Gamma\rangle$ (for some appropriate range of concepts of higher order properties). This last result entails that a right account of the possession conditions of $\langle\Phi\rangle$ requires the attribution of reference to $\langle\Phi\rangle$.⁴⁷ Consequently our perceptual ascription of properties as categorical qualifies as a variety of reference.

This variety of reference captures what was essential to the name/bearer paradigm. For just as grasp of a criterion of identity permits a speaker to distinguish an object from other objects, a

⁴⁷ My solution does not require that we classify properties as objects (and thus isn’t analogous to insisting that some putative predicates are singular terms); but the fact that you have to bring Φ under Γ provides a ground for thinking that something like the name/bearer paradigm is at work. Compare Dummett (1973/1981, p. 78): ‘What determines whether a word is to be classified as a singular term or a predicate is whether or not it occurs in other contexts in which predicates are attached to it. When its occurrence in such contexts is readily dispensable, as in the case of ‘wisdom’, it is therefore quite reasonable to regard it as only a spurious singular term, a stylistic variant on the corresponding predicate.’ Dummett later suggests that the dispensability of higher order predication—as in the case of ‘wisdom’—tracks the richness of our vocabulary (and in particular whether we have an rich array of expressions for properties of that sort). In contrast, my central move has been to show that we must, at a quite fundamental level of thought, recognise properties as having higher order properties.

subject who perceptually ascribes a property as categorical can distinguish the property by how, in conjunction with other properties, it contributes to change (that is, by the sort of difference the property makes to objects that have it). Grasp of a criterion of identity for an object will still involve something absent in the property case—namely sensitivity to the persistence conditions of the object (e.g. to grasp a criterion of identity for a chair, I must have some sense of which changes the chair could survive). But I've shown that perceptual ascription of observable properties as categorical involves something just as demanding: someone who perceptually ascribes a particular shape as categorical, for instance, must know how this property fits relative to our primitive understanding of basic causal transactions.

This reply to Dummett's challenge does not straightforwardly generalise to other varieties of property ascription. The explanatory role that the BRIDGE PREMISS affords our capacity to perceptually ascribe properties as categorical turns out to be quite special. Indeed, the *fundamental* character of this explanatory role helps to make my solution available.

I think the fact that my solution fails to generalise ultimately counts as a virtue rather than a vice. For this restricted solution respects intuitive limits on reference. Dummett was not crazy to question Frege's insistence that predicative expressions refer; his challenge has an intuitive foundation. Predicative expressions and concepts are primarily devices of classification: we use them to group like with like, or to separate some things from others. It is precisely this classificatory role that for the most part makes predicative thought and talk cheap. The introduction of new predicative concepts or expressions usually requires nothing more than a minimal set of standards for applying the concept or expression. Yet at the same time not all predicative thought and talk is cheap. Perception constrains the stock of perception-based predicates we're willing to introduce. The activity of introducing these predicates ceases to be unconstrained when I find myself in perceptual contact with ways things are. My solution to Dummett's challenge nicely explains this intuitive asymmetry between perception-based property ascription and other kinds of property

ascription: only the former can count as ascription of a property as categorical. And it is our capacity to ascribe properties as categorical that must underwrite our capacity to ascribe causation, a role which in turn makes this way of ascribing properties a variety of reference.

Bibliography

- Anscombe, G. E. (1971). *Causality and Determination*. Cambridge: Cambridge University Press.
- Antony, L. (2011). The Openness of Illusions. *Philosophical Issues*, 21: 25-44.
- Armstrong, D. (1955). Illusions of Sense. *Australasian Journal of Philosophy*, 33 (2): 88-106.
- Armstrong, D. (1997). *A World of States of Affairs*. Cambridge: Cambridge University Press.
- Austin, J. L. (1962). *Sense and Sensibilia*. (J. L. Warnock, Ed.) Oxford: Oxford University Press.
- Ayers, M. (1974). Individuals without Sortals. *Canadian Journal of Philosophy*, 4 (1): 113-148.
- Ayers, M. (1991). Substance: Prolegomena to a Realist Theory of Identity. *Journal of Philosophy*, 88 (2): 69-90.
- Ayers, M. (1997). Is "Physical Object" a Sortal Concept? A Reply to Xu. *Mind and Language*, 12 (3-4): 393-405.
- Barwise, J., & Cooper, R. (1981). Generalized Quantifiers and Natural Language. *Linguistics and Philosophy*, 4: 159-219.
- Beaney, M. (Ed.). (1997). *The Frege Reader*. Oxford: Blackwell.
- Beck, J. (2012). The Generality Constraint and the Structure of Thought. *Mind*, 121 (483): 563-600.
- Black, M., & Geach, P. T. (Eds.). (1952). *Translations from the Philosophical Writings of Gottlob Frege*. Oxford: Blackwell.
- Bostock, D. (2012). *Russell's Logical Atomism*. Oxford: Oxford University Press.
- Braun, D. (1993). Empty Names. *Nous*, 27 (4): 449-469.
- Brewer, B. (1999). *Perception and Reason*. Oxford: Oxford University Press.
- Brewer, B. (2005). Do Sense Experiential States Have Conceptual Content? In M. Steup, & E. Sosa (Eds.), *Contemporary Debates In Epistemology*. Oxford: Blackwell.

- Brewer, B. (2011). *Perception and Its Objects*. Oxford: Oxford University Press.
- Burge, T. (1973). Reference and Proper Names. *Journal of Philosophy*, 70 (14): 425-439.
- Burge, T. (1993). Content Preservation. *Philosophical Review*, 102 (4): 457-488.
- Burnyeat, M. (1990). *The Theaetetus of Plato*. Indianapolis: Hackett.
- Button, T. (2013). *Limits of Realism*. Oxford: Oxford University Press.
- Byrne, A. (2000). Intentionalism Defended. *Philosophical Review*, 110 (2): 199-240.
- Byrne, A. (2009). Experience and Content. *Philosophical Quarterly*, 59 (236): 429-451.
- Camp, E. (2004). The Generality Constraint and Categorial Restrictions. *Philosophical Quarterly*, 54 (215): 209-230.
- Camp, E. (2009). Putting Thoughts To Work: Concepts, Systematicity, and Stimulus-Independence. *Philosophy and Phenomenological Research*, 78 (2): 275-311.
- Campbell, J. (1986). Conceptual Structure. In C. Travis (Ed.), *Meaning and Interpretation*. Oxford: Blackwell.
- Campbell, J. (1993). A Simple View of Colour. In J. Haldane, & C. Wright (Eds.), *Reality: Representation and Projection*. Oxford: Oxford University Press.
- Campbell, J. (2002). *Reference and Consciousness*. Oxford: Oxford University Press.
- Campbell, J. (2006a). Does Visual Reference Depend on Sortal Classification? Reply to Clark. *Philosophical Studies*, 127 (2): 221-237.
- Campbell, J. (2006b). Manipulating Colour: Pounding an Almond. In T. S. Gendler, & J. Hawthorne (Eds.), *Perceptual Experience*. Oxford: Oxford University Press.
- Campbell, J. (2009). Consciousness and Reference. In B. McLaughlin, A. Beckermann, & S. Walter (Eds.), *Oxford Handbook of Philosophy of Mind*. Oxford: Oxford University Press.
- Campbell, J., & Cassam, Q. (2014). *Berkeley's Puzzle*. Oxford: Oxford University Press.

- Carrasco, M. (2011). Visual Attention: The Past 25 Years. *Vision Research*, 51: 1484-1525.
- Chalmers, D. (2004). The Representational Character of Experience. In B. Leiter (Ed.), *The Future for Philosophy*. Oxford: Oxford University Press.
- Chalmers, D. (2006a). Perception and the Fall from Eden. In T. S. Gendler, & J. Hawthorne (Eds.), *Perceptual Experience*. Oxford: Oxford University Press.
- Chalmers, D. (2006b). The Foundations of Two-Dimensional Semantics. In M. Garcia-Carpintero, & J. Macia (Eds.), *Two-Dimensional Semantics: Foundations and Applications*. Oxford: Oxford University Press.
- Chalmers, D. (2012). *Constructing the World*. Oxford: Oxford University Press.
- Chalmers, D. (forthcoming). Three Puzzles about Spatial Experience. In A. P. Stoljar (Ed.), *Themes from Ned Block*. Oxford: Oxford University Press.
- Cohen, J. (2010). Color Relationalism and Color Phenomenology. In *Perceiving the World*. Oxford: Oxford University Press.
- Crane, T. (1988). The Waterfall Illusion. *Analysis*, 48: 142-147.
- Dasgupta, S. (2009). Individuals: An Essay In Revisionary Metaphysics. *Philosophical Studies*, 145 (35): 35-67.
- Dickie, I. (2010a). The Generality of Particular Thought. *Philosophical Quarterly*, 60 (240): 508-531.
- Dickie, I. (2010b). We Are Acquainted With Ordinary Things. In *New Essays on Singular Thought*. Oxford: Oxford University Press.
- Dickie, I. (2012). Skill Before Knowledge. *Philosophy and Phenomenological Research*, 85 (3): 737-745.
- Dickie, I. (2014). The Sortal Dependence of Demonstrative Reference. *European Journal of Philosophy*, 22 (1): 34-60.
- Dickie, I. (2015). *Fixing Reference*. Oxford: Oxford University Press.

- Dickie, I., & Rattan, G. (2010). Sense, Communication, and Rational Engagement. *Dialectica*, 64 (2): 131-151.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge: MIT Press.
- Dretske, F. (1995). *Naturalizing the Mind*. Cambridge: MIT Press.
- Dummett, M. (1973/1981). *Frege: Philosophy of Language* (2nd ed.). Cambridge: Harvard University Press.
- Dummett, M. (1981). *The Interpretation of Frege's Philosophy*. Cambridge: Harvard University Press.
- Dummett, M. (1991a). *The Logical Basis of Metaphysics*. Cambridge: Harvard University Press.
- Dummett, M. (1991b). *Frege: Philosophy of Mathematics*. Cambridge: Harvard University Press.
- Dummett, M. (1994). Testimony and Memory. In A. C. Matilal (Ed.), *Knowing From Words*. Kluwer.
- Eklund, M. (n.d.). Regress, Unity, Facts, and Propositions.
- Evans, G. (1981). Understanding Demonstratives. In H. Parret (Ed.), *Meaning and Understanding*. Oxford: Clarendon Press.
- Evans, G. (1982). *The Varieties of Reference*. (J. McDowell, Ed.) Oxford: Oxford University Press.
- Evans, G. (1985). Things Without the Mind. In *Collected Papers*. Oxford: Oxford University Press.
- Fara, D. G. (2015). Names are Predicates. *Philosophical Review*, 124 (1): 59-117.
- Field, H. (1978). Mental Representation. *Erkenntnis*, 13 (1): 9-61.
- Field, H. (1994). Deflationist Views of Meaning and Content. *Mind*, 103 (411): 249-285.
- Fine, K. (2001). The Question of Realism. *Philosophers' Imprint*, 1 (2): 1-30.
- Fish, W. (2009). *Perception, Hallucination, and Illusion*. Oxford: Oxford University Press.
- Fodor, J. (1987). *Psychosemantics*. Cambridge: MIT Press.
- Frege, G. (1879). *Begriffsschrift*. Halle.

- Frege, G. (1884/1950). *The Foundations of Arithmetic*. (J. L. Austin, Trans.) Evanston: Northwestern University Press.
- Frege, G. (1979). *Posthumous Writings*. (H. Hermes, F. Kambartel, F. Kaulbach, Eds., P. Long, & R. White, Trans.) Oxford: Blackwell.
- Funkhouser, E. (2006). The Determinable-Determinate Relation. *Nous*, 40: 548-569.
- Furth, M. (1968). Two Types of Denotation. In N. Rescher (Ed.), *Studies in Logical Theory: American Philosophical Quarterly Monograph*. Oxford: Blackwell.
- Genone, J. (2014). Appearance and Illusion. *Mind*, 123 (490): 339-376.
- Gibson, J. J. (1977). The Theory of Affordances. In R. Shaw, & J. Bransford (Eds.), *Perceiving, Acting, and Knowing: Toward An Ecological Psychology*. Erlbaum.
- Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Houghtin Mifflin.
- Goldman, A. (1979). What is Justified Belief? In G. Pappas (Ed.), *Justification and Knowledge*. Boston: D. Reidel .
- Goodman, R. (2012). Why and How Not to Be a Sortalist about Thought. *Philosophical Perspectives*, 26 (1): 77-112.
- Hardin, C. L. (1988). *Color for Philosophers*. Indianapolis: Hackett.
- Harman, G. (1990). The Intrinsic Quality of Experience. *Philosophical Perspectives*, 4: 31-52.
- Heck, R. (2000). Nonconceptual Content and the ‘Space of Reasons’. *Philosophical Review*, 109 (4): 483-523.
- Heck, R., & May, R. (2011). The Composition of Thoughts. *Nous*, 45: 126-166.
- Heim, I., & Kratzer, A. (1998). *Semantics in Generative Grammar*. Oxford: Blackwell.
- Hellie, B. (2010). An Externalist's Guide to Inner Experience. In B. Nanay (Ed.), *Perceiving the World*. Oxford: Oxford University Press.
- Hellie, B. (2011). There It Is. *Philosophical Issues*, 21: 110-164.

- Hellie, B. (2014). Love in the Time of Cholera. In B. Brogaard (Ed.), *Does Perception Have Content?* Oxford: Oxford University Press.
- Hoerl, C., McCormack, T., & Beck, S. R. (Eds.). (2011). *Understanding Counterfactuals, Understanding Causation: Issues in Philosophy and Psychology*. Oxford: Oxford University Press.
- Hume, D. (1739). *A Treatise of Human Nature* (1896 Edition ed.). (L. A. Selby-Bigge, Ed.) Oxford: Clarendon Press.
- Jackson, F. (1977). *Perception: A Representative Theory*. Cambridge: Cambridge University Press.
- Jeshion, R. (2002). Acquaintanceless De Re Belief. In J. Campbell, M. O'Rourke, & D. Shier (Eds.), *Meaning and Truth: Investigations in Philosophical Semantics*. Seven Bridges Press.
- Johnston, M. (2004). The Obscure Object of Hallucination. *Philosophical Studies*, 120: 113-183.
- Johnston, M. (2014). The Problem With The Content View. In B. Brogaard (Ed.), *Does Perception Have Content?* Oxford: Oxford University Press.
- Johnston, W. E. (1921). *Logic, Part I*. Cambridge: Cambridge University Press.
- Kalderon, M. E. (2008). Metamerism, Constancy, and Knowing Which. *Mind*, 117 (468): 935-971.
- Kalderon, M. E. (2011). Color Illusion. *Nous*, 45 (4): 751-775.
- Kaplan, D. (1989). Afterthoughts. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes From Kaplan*. Oxford: Oxford University Press.
- King, J. C. (2007). *The Nature and Structure of Content*. Oxford: Oxford University Press.
- Kripke, S. (1980). *Naming and Necessity*. Cambridge: Harvard University Press.
- Levine, J. (2003). Experience and Representation. In Q. Smith, & A. Jokic (Eds.), *Consciousness: New Philosophical Perspectives*. Oxford: Oxford University Press.
- Lewis, D. (1970). General Semantics. *Synthese*, 22 (1-2): 18-67.
- Lewis, D. (1973). Causation. *Journal of Philosophy*, 70: 556-567.

- Lewis, D. (1980a). Index, Context, and Content. In S. Kanger, & S. Ohman (Eds.), *Philosophy and Grammar*. Dordrecht: Reidel.
- Lewis, D. (1980b). Veridical Hallucination and Prosthetic Vision. *Australasian Journal of Philosophy*, 58 (3): 239-249.
- Lewis, D. (1983). New Work for a Theory of Universals. *Australasian Journal of Philosophy*, 61: 343-377.
- Lewis, D. (1984). Putnam's Paradox. *Australasian Journal of Philosophy*, 62 (3): 221-236.
- Lewis, D. (1986). *On the Plurality of Worlds*. Oxford: Blackwell.
- Lewis, D. (1994). Reduction of Mind. In S. Guttenplan (Ed.), *Companion to Philosophy of Mind* (pp. 412-431). Oxford: Blackwell.
- Lewis, D. (2000). Causation as Influence. *Journal of Philosophy*, 97 (4): 182-197.
- Liebesman, D. (2015). Predication as Ascription. *Mind*, 124 (494): 517-569.
- Logue, H. (2012). Why Naive Realism? *Proceedings of the Aristotelian Society*, 112 (2): 211-237.
- Lycan, W. (1996). *Consciousness and Experience*. Cambridge: MIT Press.
- Martin, M. (2002). The Transparency of Experience. *Mind and Language*, 4: 376-425.
- Martin, M. (2004). The Limits of Self-Awareness. *Philosophical Studies*, 120: 37-89.
- Matthen, M. (2005). *Seeing, Doing, and Knowing*. Oxford: Clarendon Press.
- Matthen, M. (2010). Color Experience: A Semantic Theory. In M. Matthen, & J. Cohen (Eds.), *Color Ontology and Color Science*. Cambridge: MIT Press.
- McDowell, J. (1994). *Mind and World*. Cambridge, MA: Harvard University Press.
- Millikan, R. G. (1989). Biosemantics. *Journal of Philosophy*, 86 (6): 281-297.
- Palmer, S. (1999). *Vision Science*. Cambridge: MIT Press.

- Pautz, A. (2006). Sensory Awareness is Not a Wide Physical Relation. *Nous*, 40: 205-240.
- Pautz, A. (2010). Why Explain Experience in Terms of Content? In B. Nanay (Ed.), *Perceiving the World*. Oxford: Oxford University Press.
- Peacocke, C. (1983). *Sense and Content*. Oxford: Clarendon Press.
- Peacocke, C. (1992). *A Study of Concepts*. Cambridge: MIT Press.
- Peacocke, C. (2008). *Truly Understood*. Oxford: Oxford University Press.
- Peacocke, C. (2011). Representing Causality. In T. McCormack, C. Hoerl, & S. Butterfill (Eds.), *Tool Use and Causal Cognition*. Oxford: Oxford University Press.
- Perry, J. (1977). Frege on Demonstratives. *Philosophical Review*, 86 (4): 474-497.
- Prosser, S. (2011). Affordances and Phenomenal Character in Spatial Perception. *Philosophical Review*, 120 (4): 475-513.
- Pryor, J. (2000). The Skeptic and the Dogmatist. *Nous*, 34 (4): 517-549.
- Putnam, H. (1981). *Reason, Truth, and History*. Cambridge: Cambridge University Press.
- Quine, W. V. (1950). Identity, Ostension, and Hypostesis. *Journal of Philosophy*, 47 (22): 621-633.
- Quine, W. V. (1960). *Word and Object*. Cambridge: MIT Press.
- Rabern, B. (2012). Against the Identification of Assertoric Content with Compositional Value. *Synthese*, 189 (1): 75-96.
- Raffman, D. (2014). *Unruly Words*. Oxford: Oxford University Press.
- Rattan, G. (2010). Intellect and Concept. *Baltic International Yearbook of Cognition, Logic, and Communication*, 5 (1): 1-61.
- Rayo, A., & Yablo, S. (2001). Nominalism Through De-Nominalization. *Nous*, 35 (1): 74-92.
- Recanati, F. (2012). *Mental Files*. Oxford: Oxford University Press.

- Reimer, M. (2002). Do Adjectives Conform to Compositionality? *Nous*, 36 (16): 183-198.
- Rieppel, M. (forthcoming). Being Something: Properties and Predicative Quantification. *Mind*.
- Roessler, J. (2011). Perceptual Causality, Counterfactuals, and Special Causal Concepts. In C. Hoerl, T. McCormack, & S. R. Beck (Eds.), *Understanding Counterfactuals and Causation: Issues in Philosophy and Psychology*. Oxford: Oxford University Press.
- Rumfitt, I. (1994). Frege's Theory of Predication: An Elaboration and Defense, with Some New Applications. *Philosophical Review*, 103 (4): 599-637.
- Rumfitt, I. (1995). Truth Wronged: Crispin Wright's Truth and Objectivity. *Ratio*, 8 (1): 100-107.
- Rumfitt, I. (2000). "Yes" and "No". *Mind*, 109 (436): 781-823.
- Russell, B. (1905). On Denoting. *Mind*, 14: 479–493.
- Russell, B. (1910). Knowledge by Acquaintance and Knowledge by Description. *Proceedings of the Aristotelian Society*, 11: 108-128.
- Russell, B. (1912). *The Problems of Philosophy*. Routledge.
- Russell, B. (1913/1984). *Theory of Knowledge: The 1913 Manuscript*. London: Allen and Unwin.
- Russell, B. (1918/1956). The Philosophy of Logical Atomism. In R. Marsh (Ed.), *Logic and Knowledge*. London: Allen and Unwin.
- Sainsbury, M. (1996). How Can We Say Something? In R. Monk, & A. Palmer (Eds.), *Bertrand Russell and the Origins of Analytical Philosophy*. Thoemmes.
- Schellenberg, S. (2010). The Particularity and Phenomenology of Perceptual Experience. *Philosophical Studies*, 149 (1): 19-48.
- Scholl, B. (2001). Objects and Attention: The State of the Art. *Cognition*, 80 (1/2): 1-46.
- Shoemaker, S. (1980). Causality and Properties. In P. v. Inwagen (Ed.), *Time and Cause*. Dordrecht: D. Reidel.
- Shoemaker, S. (1988). On What There Are. *Philosophical Topics*, 16 (1): 201-223.

- Shoemaker, S. (1994). Phenomenal Character. *Nous*, 28: 21-38.
- Siegel, S. (2010). *The Contents of Visual Experience*. Oxford: Oxford University Press.
- Siegel, S. (2014). Affordances and the Content of Perception. In B. Brogaard (Ed.), *Does Perception Have Content?* Oxford: Oxford University Press.
- Siewert, C. (1998). *The Significance of Consciousness*. Princeton: Princeton University Press.
- Silins, N. (2011). Seeing Through The 'Veil of Perception'. *Mind*, 120 (478): 329-367.
- Smithies, D. (2011). What is the Role of Consciousness in Demonstrative Thought? *Journal of Philosophy*, 108 (1): 5-34.
- Soames, S. (2010). *What Is Meaning?* Princeton: Princeton University Press.
- Soteriou, M. (2005). The Subjective View of Experience and its Objective Commitments. *Proceedings of the Aristotelian Society*, 105 (2): 177-190.
- Soteriou, M. (2013). *The Mind's Construction*. Oxford: Oxford University Press.
- Stalnaker, R. (1981). Indexical Belief. *Synthese*, 49 (1): 129-151.
- Stalnaker, R. (1984). *Inquiry*. Cambridge: MIT Press.
- Stanley, J. (1997). Rigidity and Content. In R. Heck (Ed.), *Language, Thought, and Logic: Essays in Honour of Michael Dummett*. Oxford: Oxford University Press.
- Stanley, J. (2011). *Know How*. Oxford: Oxford University Press.
- Strawson, P. F. (1959). *Individuals*. London: Methuen.
- Strawson, P. F. (1974). *Subject and Predicate in Logic and Grammar*. London: Methuen.
- Strawson, P. F. (1979). Perception and Its Objects. In G. F. MacDonald (Ed.), *Perception and Identity: Essays Presented to A. J. Ayer*. London: MacMillan.
- Strawson, P. F. (1992). *Analysis and Metaphysics*. Oxford: Oxford University Press.

- Thau, M. (2002). *Consciousness and Cognition*. Oxford: Oxford University Press.
- Thompson, B. (2010). The Spatial Content of Experience. *Philosophy and Phenomenological Research*, 81 (1): 146-184.
- Travis, C. (1994). On Constraints of Generality. *Proceedings of the Aristotelian Society*, 94: 165-188.
- Travis, C. (2004). The Silence of the Senses. *Mind*, 113 (449): 57-94.
- Tye, M. (1995). *Ten Problems of Consciousness*. Cambridge: MIT Press.
- Tye, M. (2000). *Consciousness, Color, and Content*. Cambridge: MIT Press.
- Wedgwood, R. (2007). *The Nature of Normativity*. Oxford: Oxford University Press.
- Wiggins, D. (1997). Sortal Concepts: A Reply to Xu. *Mind and Language*, 12 (3&4): 413–421.
- Williams, D. C. (1953). The Elements of Being. *Review of Metaphysics*, 7 (2), 3-18: 171-192.
- Williams, J. R. (2007). Eligibility and Inscrutability. *Philosophical Review*, 116 (3): 361-399.
- Wittgenstein, L. (1921). *Tractatus Logico-Philosophicus*. (D. Pears, & B. McGuinness, Eds.) London: Routledge.
- Wright, C. (1975). On the Coherence of Vague Predicates. *Synthese*, 30 (3/4): 325-365.
- Wright, C. (1983). *Frege's Conception of Numbers as Objects*. Aberdeen: Aberdeen University Press.
- Wright, C. (1994). *Truth and Objectivity*. Cambridge: Harvard University Press.
- Wright, C. (2007). On Quantifying Into Predicate Position. In M. Leng, & A. Paseau (Eds.), *Mathematical Knowledge*. Oxford: Oxford University Press.
- Wu, W. (2014). *Attention*. London: Routledge.
- Yalcin, S. (2007). Epistemic Modals. *Mind*, 116 (464): 983-1026.

