

Perception and Inference

by

Thomas J. Breed

B.A., Philosophy, University of Warwick, 2010

M.A., Language and Cognition, King's College London, 2013

Submitted to the Graduate Faculty of

the Kenneth P. Dietrich School of Arts and Sciences in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2020

UNIVERSITY OF PITTSBURGH

THE KENNETH P. DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Thomas J. Breed

It was defended on

July 28th 2020

and approved by

John McDowell, Distinguished University Professor of Philosophy

Robert Brandom, Distinguished Professor of Philosophy

Anil Gupta, Alan Ross Anderson Distinguished Professor of Philosophy

Wayne Wu, Associate Professor, Carnegie Mellon University

Dissertation Director: John McDowell, Distinguished University Professor of Philosophy

Copyright © by Thomas J. Breed

2020

Perception and Inference

Thomas J. Breed, PhD

University of Pittsburgh, 2020

This project explores the connections between normativity, perception, and conceptual content, and develops a view on which perception has inferentially-articulated conceptual content that is grounded in the way perceptual experiences alter the normative statuses of perceivers.

In chapter two I address a criticism developed by Tyler Burge, who argues that views that require perceivers to have sophisticated capacities, such as conceptual or inferential capacities, in order to have representationally contentful perceptual states conflict with our best understanding of vision science. I argue that even if, as Burge claims, non-rational creatures have representational perceptual states that do not depend in any way on such capacities, nonetheless the perceptual states of *rational* creatures must be integrated with their conceptually contentful mental states and must be capable of serving as reasons for belief, given the role these states must play in the explanation of the representational perspectives of rational creatures.

In chapter three I develop and defend a specific version of conceptualism on which the conceptual content of perception is grounded in its normative significance. Specifically, I claim, by analogy with inferentialist metasemantic accounts of meaning, that perception has conceptual content by virtue of the way that it warrants perceivers in applying concepts and entitles them to claims. I also argue, against Robert Brandom, that giving perception this kind of role is required in order to give an adequate semantics for natural languages. I then apply this view to provide an explanation of the object-directedness of perception.

In chapter four I pair this inferentialist account of perceptual content with a relationalist approach to the metaphysics and phenomenology of perception. Most relationalists argue that their view is incompatible with the claim that perception has content, but I argue that the way the inferentialist account grounds the content of perception in normative features of experience makes it more compatible with relationalist approaches than many other similar views. I then respond to a specific problem raised for the content view by Charles Travis, and provide an explanation of the relationship between the phenomenology and content of perception on the basis of my response to this problem.

Table of Contents

Preface	x
1.0 Introduction	1
1.1 Chapter Two: Tyler Burge on Objectivity	2
1.2 Chapter Three: Inference and Perception	6
1.3 Chapter Four: The Relational Aspect of Perception	10
2.0 Objectivity and Rational Perception	14
2.1 Introduction	14
2.2 Tyler Burge on Objectivity	18
2.2.1 Second-Family Individual Representationalism	18
2.2.2 Objectivity and Objectification	24
2.2.3 Transformative Conceptions of Rationality	28
2.2.4 Burge’s Blind-Spot	32
2.3 Perception and Explanation	38
2.3.1 Burge’s Two Arguments	38
2.3.2 Burge’s Transformative Conception of Objectivity	42
2.4 Perception and Rationality	46
2.4.1 Individuals, Persons, and Reasons	46
2.4.2 Rational Perspectives and Rational Perception	50
2.4.3 Some Responses	54
2.5 Conclusion	60
3.0 An Inferentialist Account of Perception	64

3.1	Introduction	64
3.2	Inferentialism	68
3.2.1	Inferentialist Semantics	68
3.2.2	What <i>Are</i> Premises and Conclusions, Anyway?	76
3.3	From Language to Perception	80
3.3.1	Brandomian Inferentialism is a Hyperinferentialism	80
3.3.2	Strong Inferentialism and the Space of Reasons	84
3.4	A Foundational Inferentialist Account of Perceptual Content	86
3.4.1	Inferentialist Metasemantics	86
3.4.2	A Foundational Account of Perceptual Content	91
3.5	Givenness and the Inferentialist View of Perception	102
3.5.1	Object-Awareness and The Myth of the Given	102
3.5.2	Intentionality and Inference	104
3.5.3	Perception and Objects in Language Acquisition	110
3.6	Conclusion	116
4.0	Perception as a Normatively Significant Relation to Mind-Independent Objects	120
4.1	Introduction	120
4.2	Naïve Realism and Representationalism	123
4.2.1	Representationalism in the Philosophy of Perception	123
4.2.2	Inferentialism and Representationalism	127
4.2.3	The Three Central Claims of Naïve Realism	132
4.3	Inferentialism for Relationalists	139
4.3.1	Naïve Realism and the Content View	139

4.3.2	The Argument from Normative Significance	146
4.4	Contentful Visual Awareness	152
4.4.1	Travis’s Argument	152
4.4.2	Visual Looks and Visual Context	158
4.5	The Visual and the Conceptual	164
4.5.1	Appearances and the Content of Perception	164
4.5.2	Inferential Role and Inferential Significance	170
4.5.3	The Inferential Content of Perception	174
4.6	Conclusion	178
5.0	Bibliography	181

List of Figures

1	The Normative Pragmatic Inferentialist Metasemantics	90
2	The Normative Pragmatic Account of Perceptual Content	97
3	The Inferential Role of a Sentence	171
4	Same Inferential Role, Different Inferential Content	173
5	The Inferential Content of Daylight and Dive Bar	175

Preface

“The subject of genuine perceptual beliefs is... responding to the visible presence of red things by making a potential move in a game of giving and asking for reasons: applying a concept.”

—Robert Brandom, *Articulating Reasons*, p.108

“It’s all about the game and how you play it...
All about your debt and if you can pay it.”

—Motörhead, ‘The Game’

Every dissertation is the product of the ideas, influence, and support of a huge community of people, but this is especially true of the present one. As a first-generation graduate student, a lot had to happen for me to even make it to graduate school, let alone to write a dissertation. In particular, I want to thank Rebecca Jinks, Guy Longworth, and Miguel de Beistegui, for having faith in me as a student. Without their encouragement I’m not sure I would have applied to graduate school at all.

The dissertation obviously owes a great deal to my graduate committee, and especially to my advisor, John McDowell, who has often been a more persuasive expositor of the views of the people with whom I express disagreement in this dissertation than they are themselves. If I have succeeded in engaging charitably with my interlocutors at all, this is in no small part thanks to his influence. Aside from their input on the dissertation itself, many parts of this project began with the philosophy of perception class I took with Anil Gupta, and directed studies that I did with John McDowell and Bob Brandom. Much of this dissertation is in fact an attempt to frame insights I’ve gleaned from John and Anil using conceptual

tools made available by Bob. The result is almost certainly something that satisfies nobody except (perhaps) me, so for this I owe all three both my thanks and my apologies.

My sense is that my process of writing the dissertation has been somewhat atypical. When I first arrived at Pitt and asked about how people decide what to write for their dissertation, an older student joked that the official process is to wander off into the woods for a few months until you return with a fully-formed dissertation topic, and only then should you meet regularly with your committee. My experience has been, if anything, the opposite of this: Meetings with my committee were indispensable for conceiving of the project and its basic outlines, and then I wandered off into the woods for a couple of years only to return with a fully-formed dissertation (as well as an unexpected AOC in Chinese Philosophy).

One consequence of this slightly unusual writing process is that feedback that I've received from people outside of my committee has had more of an influence on the dissertation than is perhaps common. In particular, I want to thank attendees of the dissertation seminar and of various WiPs throughout the years, and to audiences at the CUA conference on cognition and the MLAG conference in Porto, especially Alva Noë and Charles Travis. The semester I spent at KCL, and the discussions I had with Bill Brewer and Matt Soteriou, also had a significant impact on the dissertation, especially on chapter 4.

A special mention goes to all of the attendees of the NEH-funded Summer Seminar 'Responses to Empiricism in Kant, Hegel, and Sellars', and to Bill deVries and Jim O'Shea for organising the event.¹ The four weeks spent at this seminar were some of the most intellectually stimulating weeks of my academic life, and affected the way I thought about this project significantly. Any hints of influence from Kant, Hegel, Sellars, and Davidson in

¹I should also acknowledge here my gratitude both to the National Endowment for the Humanities for providing a stipend to cover the costs of attending this seminar, and to the Rescher Dissertation Fellowship, which provided crucial financial support during the completion of this dissertation.

these chapters owe at least in part to the countless hours of discussion with them all during my time in New Hampshire.

More generally, the innumerable conversations I have had with members of the Pitt philosophy graduate community have helped to shape this dissertation, and to shape me both as a philosopher and as a person, more than I could hope to express here. There are too many people to name everyone individually, but I want to specifically acknowledge my fellow office-mates and Friends of the German Tradition Jack Samuel, Alnica Visser, Aaron Segal, Dan Kaplan, and Max Tegtmeyer, as well as Liam Bright, Alison Springle, David de Bruijn, Michael Lang, and Travis McKenna, who have provided not just intellectual engagement but the kind of friendship without which the process of writing a dissertation would have been much less tolerable.

In addition to these general debts, several people gave feedback that led me to make specific changes to the dissertation. Chapter 2 was initially just a single section in an early draft of a chapter, but Matt Soteriou and Wayne Wu both helped me to realise that I needed to engage with Burge more seriously. Matt's comments on an early chapter draft also led to me changing the way I dealt with appearances, which significantly impacted §4.5.3. In addition, Pablo Zendejas Medina pushed me to think about the characterisation I gave of rational perspectives in §2.4.1, which eventually led to me changing this section significantly.

Finally, more than anyone else I owe thanks to my parents. Without their hard work I couldn't hope to be even starting a PhD, and without their continual encouragement I couldn't possibly hope to be finishing one. I couldn't ask for more supportive parents, and I dedicate this dissertation to them both.

The last time I visited my family in the UK my dad sat me down and, in a serious tone of voice, told me that he and my brother needed to ask me a question. "We found this book

by Wilfrid Sellars lying on the floor”, he said, “and we want to know what in the world ‘empiricism’ means”. I’m still not exactly sure I know myself, but hopefully this dissertation provides a few clues.

1.0 Introduction

“We are therefore said to be justified by faith, because faith is the beginning of human salvation, the foundation and root of all justification, *without which it is impossible to please God* (Heb 11:6)... and we are therefore said to be justified gratuitously, because none of those things that precede justification, whether faith or works, merit the grace of justification. For, *if by grace, it is not now by works, otherwise*, as the Apostle says, *grace is no more grace* (Rom 11:6).”

—Council of Trent, *Decree Concerning Justification*, Ch.XIII

Each of us, simply by virtue of being the kinds of creatures we are, is a bearer of normative statuses. We are creatures whose behaviours and activities have a certain kind of significance—we do not merely act, but we act in ways that we *may* or *must*. Some normative statuses we bear only because of things that we have done—things we are *responsible* for. But other normative statuses are unearned. Things *happen* to us, things that we are not responsible for and did not bring about by our own volition, but which nonetheless give us *permission* to do or say things we would otherwise not be entitled to do or say.

Only creatures who are capable of *doing* things that have the first kind of normative significance are capable of *undergoing* events that have this second kind of normative significance. One cannot meaningfully be said to have permission to do things that one is not even capable of doing. But *if* one is a creature capable of bearing normative statuses at all, there must be some that one did not earn; statuses that were bestowed upon one not on the basis of what one did but merely by virtue of the fact that something happened to one.

This dissertation is, at its heart, about how it is that this is possible and what the significance of this fact is for one small but important area of our lives—perception. I argue

that perception bestows upon us normative statuses without us having to do anything to earn them, and that this has important implications for how we think about perception. Specifically, I argue that the distinctive *kind* of normative significance that perception has is that it *warrants us in applying concepts* and *entitles us to claims*, and I argue that by virtue of doing this perception has *conceptual content*.

In the rest of this introduction I want to briefly sketch the argument I give for this view in the three central chapters of this dissertation.

1.1 Chapter Two: Tyler Burge on Objectivity

In chapter 2, I clear the ground for my account by addressing a problem faced by any view that claims to provide an explanation of the content of perception in terms that make essential appeal to our ability to apply concepts, or by reference to our capacity to reflect on reasons. The problem, in short, is that there is good reason to think that the fact that perception has content is what makes it *objective*. The difference between mere sensory discrimination—the sort of awareness had even by simple creatures such as earthworms and amoebae—and genuinely *objective* perception, according to this line of thought, is that the latter is *representational*. Genuinely objective states do not merely carry information about the environment, but have content that determines the way the environment would have to be for that experience to be accurate.

The problem faced by a view that claims to ground the content of perception in sophisticated abilities such as one's ability to reflect on reasons is precisely that these abilities *are* sophisticated. Most creatures—not just earthworms and amoebae, but cats and frogs, and even infant humans—do not have these kinds of capacities. So if these capacities are required

for objective perception *tout court* then it seems like this account denies that the majority of creatures are capable of objectively perceiving. Creatures other than adult humans might be capable of sensorily discriminating features of their environment, but they could not have genuinely objective perceptual states; states that represent the environment as being certain ways.

In *Origins of Objectivity*, Tyler Burge argues that although this kind of view is incredibly widespread in philosophy, it is false. It is false, Burge claims, because many creatures other than humans are capable of objectively perceiving. Philosophers who claim that objective perception requires sophisticated intellectual capacities underestimate the sophistication of the visual system itself, according to Burge. The visual systems of a wide variety of creatures are capable of processing the information registered by sensory receptors in such a way as to determine what the distal causes of those patterns of stimulation likely are. If the creature is in a state that is the result of this kind of process, Burge claims, that is sufficient for that state to objectively represent the creature's environment as being some way, and hence for its experience to have content.

Most of the responses to Burge's argument have focused on the question of whether these sorts of processes that recover information about distal causes really are necessary for perception. The argument that I give in chapter 2 focuses on a different part of Burge's view. Even if Burge is right that for a sensory state to objectively represent it is *necessary* for that state to be generated by an objectifying process, Burge himself acknowledges that this is not *sufficient*, for the state must convey *to the creature* how things are in its environment. The state, that is, must be attributable to it as an individual rather than merely to some of its subsystems, and it must therefore be capable of featuring in explanations of that creature's behaviour and of its representational perspective.

I argue that given that rational perceivers are the bearers of conceptually contentful states such as beliefs, if perceptual states are to be attributable to rational individuals *as such*, those states must bear rational relations to the conceptually contentful states of that individual, and must be capable of serving as a reason for belief. I lay out this argument in detail in this chapter, but I focus almost exclusively on Burge’s view, trying to engage with it in detail, and I motivate my view by reference to features distinctive to Burge’s own account. Here, therefore, I want to highlight an analogy with another approach to perception that I think might be instructive as a way of drawing out the broader motivations for this kind of response.

The analogy is with Ruth Millikan’s teleosemantic account of perception. Teleosemantic accounts are views that ground the content of perception in biological notions of proper function. Those who endorse teleological views of this kind typically try to account for representational accuracy in terms of a state’s fulfilling its proper function and to account for representational error in terms of its *failing* to fulfil its proper function.

What is distinctive about Millikan’s account, as opposed to some other versions of teleosemantics, is that she focuses on the norms governing systems that *consume* representations—the systems that treat certain states *as* giving them information about the environment—rather than those that *produce* the representations. Millikan notes that once we shift our perspective to focus on what is required for something to be *treated* as a representation by the system *consuming* it, features of the system that *produce* the representation drop out of consideration as irrelevant except insofar as they are capable of being recognised by the consumer *as* relevant:

“If we know what constitutes the consumer’s *taking* a sign to indicate *p*, what *q*, what *r*, etc., then, granted that the consumer’s takings are in some way systematically derived from the structures of the signs so taken, we can construct a semantics for the consumer’s

language. Anything the signs may indicate qua natural signs or natural information carriers then drops out as entirely irrelevant; the representation-producing side of the system had better pay undivided attention to the language of its consumer.” (Millikan, 1989, p. 286)

In effect, the argument I give in chapter 2 aims to show something similar is true if we think of rational creatures as the consumers of states produced by objectifying visual capacities. It may be necessary for perception to be objective that perceptual states are produced by such processes, but any contribution these processes make to the objectivity of perceptual states must be capable of being recognised by the consumer. Burge argues, in effect, that for a consumer to be said to recognise the way things are represented to it as being it is sufficient that its perceptual states feature in explanations of its behaviour and its perspective. I argue that given that we are rational creatures, perceptual states have to be available for *rational* responses, and that in particular perceptual states feature in the explanation of the distinctively *rational* features of the individual’s perspective.

Thus, whilst there must be *some* sense in which the rational creature’s responses are dependent on the features of their perceptual states *qua* outputs of the sorts of processes Burge discusses, the form of recognition that is distinctive of us as rational creatures is *rational*—we are capable of telling whether one thing is a *reason* for another, or whether one thing we think is *incompatible* with another. The language we speak—to which, Millikan claims, the representation-producing side of the system had better pay undivided attention—is a language of *rational norms*. So if perception is genuinely attributable to *us*—is something that *we* are capable of treating as a representation—we must be capable of recognising how perceptual states bear on what we are to think. Or so I will argue.

1.2 Chapter Three: Inference and Perception

In chapter 3 I attempt to show that if perception has the kind of normative significance that I claimed, in the first chapter, it must have if it is to be objective, then such states do not merely represent, but they have *conceptual content*. The framework within which I provide this account is that of *inferentialism*. The central feature of the inferentialist approach to language relevant for understanding its bearing on perception is a commitment to *normative pragmatism* about attributions of content:

“Semantics must answer to pragmatics. The theoretical point of attributing semantic content to intentional states, attitudes, and performances is to determine the pragmatic significance of their occurrence in various contexts. This means settling how linguistic expressions of these contents are properly or correctly used, under what circumstances it is appropriate to acquire states or attitudes with those contents, and how one then ought or is obliged to go on to behave.” (Brandom, 1994a, p. 83)

The point of attributing content to an assertion, on this account, is to characterise the *normative* significance of that event. The meaning of a sentence, on this view, is its *normative inferential role*—a feature of the proper or correct use of the sentence grounded in the way that assertions of that sentence alter one’s normative statuses.

To illustrate: If Benedict says that he has a white cat then Benedict has acquired certain *commitments* he did not have before—namely, commitment to the *claim* that he has a white cat. The conceptual dimension of such commitments, on the inferentialist account, consists in their *inferential articulation*: If Benedict undertakes commitment to the claim that he has a white cat he *ipso facto* acquires commitment to the claim that he has *a* cat, that he has a pet, that there is at least one white cat, and so on.

However, linguistic commitments, like other commitments, also bring with them certain responsibilities. If I overhear Benedict telling someone else that his cat is black, or if I visit

his house and only see a dog lounging on the sofa, I might challenge him to justify what he said; to demonstrate his *entitlement* to the claim that he has a white cat. A common way to demonstrate one's entitlement is by giving a *reason*—that is, by making an assertion one's entitlement to which entitles one to the claim in question.

In this chapter I argue that perception is the ultimate source of such entitlements; that by being visually aware of something a perceiver acquires entitlement to certain claims without actually needing to undertake commitment to them. There are two implications of this fact. Firstly, if experiences serve as reasons for belief in this way, this means that inferential connections to experience partly determine the meaning of sentences themselves. One of the features of use that, on the inferentialist account, are constitutive of the meaning of a sentence are the circumstances in which it is appropriate to undertake commitment to that sentence. If perception entitles one to claims then perceptual circumstances of application themselves partly determine the meaning of sentences: What 'red' means is partly determined by the fact that I am entitled to apply the concept red to red looking things.

There is a temptation to deny that there is any need for this sort of unearned, gratuitous justification—a kind of Pelagian insistence that we don't need any help; that any positive status we might have we must have earned ourselves, on our own merit. I argue in this chapter that this is a mistake, and that in particular if we insist that the only normatively significant events are things that we *do*, such as asserting or judging, and we deny that at least *some* of these events only have the significance they do by virtue of events that we merely *undergo*, such as experiencing or perceiving, we cannot make sense of the empirical contentfulness of language at all.

The second implication of this fact is of more relevance to the philosophy of perception than to the philosophy of language. If perception can serve as a reason for belief by virtue

of the way it entitles perceivers to claims, this has implications not just for the content of language but also for the content of perception itself. By entitling perceivers to claims, perception has a normative significance akin to that of an assertion. Specifically, given that perception entitles perceivers to *claims*, experiences, like assertions, have conceptual content: If Benedict *sees* a white cat his experience might entitle him to the claim that there is a white cat in front of him. This has implications for what sorts of *other* things he is entitled to. He might not see the cat's claws, but if his experience entitles him to the claim that there is a white cat in front of him it will *ipso facto* entitle him to the claim that the creature in front of him has retractable claws. Thus, an implication of this account is that experiences have *inferentially articulated content*.

Given that the features in which the content of perception are grounded are its normative features, however, there is a question whether this account can tell us anything interesting about perception itself. Normative statuses are not *physical properties*, so if to characterise perception as having content is fundamentally just to characterise it as having a certain normative significance, such an account of content will not be able to explain, for example, why it is that perception has phenomenology, which is a common aim of many views that attribute content to perception. I thus give one example in this chapter of the kind of light that I think this sort of view can shed on our understanding of perception. Specifically, I explain what it means for perception to present a perceiver with an object *as* an object in terms of the distinctive kind of *normative inferential significance* that awareness of objects *qua* objects has.

If these entitlements are ultimately statuses that have relevance for the 'game of giving and asking for reasons'—if experience gives us justifications that we can articulate and cite in discussion with one another, say to vindicate our entitlement to a claim in response to

a challenge—they must be grounded in features of perception that people other than the perceiver herself are also capable of recognising. If the normative significance of perception were grounded in my awareness of fundamentally private features of experience inaccessible to other speakers, it is hard to see how perception could play this role. Amongst other things, it is hard to see how I could possibly *check* whether your experience really did entitle you to this claim or whether you merely *thought* that it did. This suggests that the features of perception relevant to its normative significance are features of the environment of which a perceiver is aware.

I think that this claim is correct, but in the last section of this chapter I provide a slightly different motivation for thinking that the normatively relevant features of perception are features of mind-independent objects. Specifically, I provide an explanation of how it is that perceivers come to be able to have perceptual entitlements at all—namely, by virtue of being taught a language and, in particular, by being taught how to recognise what sorts of linguistic entitlements their experiences confer upon them. If this is possible, however, teachers of the language must be able to recognise the situations in which the learners of the language have these entitlements (or, if they are particularly young students, at least *would* have these entitlements were they mature users of the language). I argue that this provides a motivation for thinking that the normatively relevant features of perception are actually just features of the perceived world, made perceptually available to both teacher and student by their experiences. In the next chapter I try to spell out in detail what an account of perception based on this suggestion would look like.

1.3 Chapter Four: The Relational Aspect of Perception

In the final chapter, chapter 4, I articulate and defend a hybrid view of visual perception according to which perception relates perceivers to mind-independent entities and I argue that, precisely by virtue of relating perceivers to these entities in the distinctively *visual* way that they do, perceptual states have conceptual content.

I begin by laying out the central claims of *naïve realism*, which, roughly, is the view that perception visually relates perceivers to mind-independent entities, and that the way things visually appear to a perceiver is determined by the visible features of the mind-independent entities to which she is related. Most naïve realists reject the claim that perception has content, and so in the first part of this chapter I address this concern, arguing that there is fact no conflict between the view that perception has content *as such* and the claim that perception fundamentally consists in a relation to mind-independent objects.

Specifically, I argue that although there are some insoluble conflicts between naïve realism and *some* versions of the content view, these conflicts are a consequence of the claim that perception has content as such but rather are a consequence of the particular kinds of explanations many philosophers give of *how* perception has content; the features of perception by virtue of which it has content. I argue that this is a result of the naturalistic motivations most such philosophers have for developing a representational theory of perception. As I noted in the previous section, many philosophers of perception explain perception in terms of its content because they hope to be able to provide an explanation of the *phenomenology* of perception, usually in terms of naturalistically acceptable ideas such as information transmission. This leads many such philosophers to think of perception as a state that merely

happens to be *caused* by entities, rather than as a state that fundamentally *relates* perceivers to such entities, which is what generates the conflict with naïve realism.

By contrast, I argue that not only is there is no conflict between the central metaphysical and phenomenological claims made by naïve realists and the view of perception that I gave in the previous chapter but that, in fact, if perception is to play the role in the explanation of the epistemology of perception that most naïve realists want to give it, it *must* have content. Perception, on this view, is a *normatively significant* relation to mind-independent objects, and by virtue of the kind of normative significance it has, it thereby has conceptual content.

I then address an objection to this kind of view; namely, I address an argument by Charles Travis that purports to show that even if content is somehow *associated* with perception, there is no way for the perceiver to be able to *recognise* what the content of her experience is on the basis of purely visual features of her experience alone. This threatens to undermine the claim that perception has content altogether, since, Travis argues, it is a requirement of this kind of account that the content of perception be conveyed to the perceiver by her experience, and if it is to do this then there must be some features of perception itself by virtue of which the perceiver is able to recognise what its content is.

I argue that in order to understand how it is that perceivers recognise what the content of their experiences is, we have to pay attention to the way that the visual features of experience hang together; to the distinctive kind of visual *unity* of the visual field as a whole. If we do, we can see how it is that by undergoing an experience in which things visually look a certain way, a perceiver is *ipso facto* in a position to tell what the content of her experience is.

In the last part of this chapter I consider the implications of this kind of hybrid view for our characterisation of the content of perception. Specifically, I argue that the features of experience that allow a perceiver to *recognise* what the content of her experience is—namely,

how things visually look to her—are the very same features that actually determine what the content of her experience is. I begin this final part of the chapter by motivating this claim. Specifically, I note that if a perceiver’s response to the way things visually look is supposed to be *rational* then we have to acknowledge that the way things look itself, and not merely the perceiver’s conceptual responses to her experience, can be conceptually articulated.

To give an account of the content of appearance claims, however, we need to bring in more of the inferentialist semantics than I outlined in the previous chapter. Specifically, we need to acknowledge two crucial features of the kinds of material inferences that, on the inferentialist account, constitute the meaning of a sentence. Firstly, what follows from a claim often depends on what other assumptions one relies on in drawing out its inferential consequences. We thus need to acknowledge an essential role for *multi-premise* inferences in our inferential semantics. Secondly, many inferences are only *defeasibly* good: There is a good inference from “Pangur is a cat” to “Pangur has fur”, but this inference can be turned into a bad one simply by adding further premises. Thus, there is *not* a good inference from “Pangur has fur” *and* “Pangur is a Sphynx” to “Pangur has fur”.

Once I have explained this part of the inferentialist semantics, I give my account of the content of visual appearances, and explain how this allows us to explain how the inferential content of perception depends on the way that the various things in a perceiver’s visual field look, and on the way that these different appearances relate to one another. Specifically, I argue that the different parts of a perceiver’s visual field stand in relation to one another as the premises of a multi-premise inference stand in relation to one another. What the content of a perceiver’s experience is depends not just on the sensory qualities present in her visual field, but on the way that these sensory qualities relate to and bear on one another. For this reason, whether an experience in which something looks red to one actually gives one reason

to believe that there is something red in front of one depends on what else one has visual awareness of.

2.0 Objectivity and Rational Perception

“There exists another power, not only that by which I give life to my body but also that by which I enable its senses to perceive. The Lord made this for me... providing the eye to see and the ear to hear, and each of the other senses in turn to be in its proper place and carry out its proper function. I who act through these diverse functions am one mind. I will also rise above this power. For this also is possessed by the horse and the mule.”

—St. Augustine, *Confessions*, X.vii

2.1 Introduction

The inferentialist account of perception I develop in subsequent chapters takes the role of perception in entitling perceivers to beliefs, and in reason-giving relations more generally, to be essential to perception’s having content (ch.3), and to the kind of awareness perception affords a perceiver of mind-independent objects (ch.4). This theory is thus part of a broader class of views I will call *intellectualist* theories of perception—theories that claim that rational capacities, such as the capacity to base beliefs on perception, and relations to rational states, such as propositional attitudes, are constitutively part of the perception of at least *rational* language-using creatures such as adult humans.¹

In this chapter I respond to what I take to be one of the most sophisticated, and certainly one of the most influential, objections to intellectualist theories of perception; namely, that

¹The use of the term ‘rational’ should not be taken to imply that the creature is *reasonable*, but *only* that it is a creature that has the ability to reflect on reasons. Whether or not it exercises these capacities *well* is another matter. I will also assume in what follows that rationality in this sense requires having a language, and so will drop the qualification ‘language-using’. There is obviously a sense in which non-language-using creatures have a certain kind of rationality, so I am using this term in a slightly more narrow sense than it is sometimes used.

provided by Tyler Burge in *Origins of Objectivity* (2010).² The target of Burge’s argument in this book is a class of views that endorse what he calls *Compensatory Individual Representationalism*, or ‘individual representationalism’ for short. Most such philosophers simply assume—out of ignorance of the science, Burge claims (Burge, 2010, p. 266)—that left to their own devices, visual systems are not capable of producing genuinely objective states. Individual representationalists thus hold that only creatures that are capable of representing some of the constitutive conditions for objective representation are capable of objectively representing *tout court*—the individual has to *compensate* for the shortcomings of the visual system.³

Thus, for example, according to some empiricist versions of individual representationalism, in order to objectively represent one must be able to think of objects as the *causes* of one’s sensations. That there are objects that *are* the causes of one’s sensations is, arguably, a constitutive condition of objective representation. But according to this version of individual representationalism, it is not enough that objects *in fact* cause the perceiver’s sensations. For the perceiver to be capable of genuinely objective representation, she must be able to think of such objects *as* the causes of her sensations. Thus, on this view, objective representation depends on the capacity to represent at least one of the constitutive conditions of objective representation.

²Aside from *Origins of Objectivity*, see especially Burge, 2003, 2005, 2009, 2014. With a few exceptions—normally in cases where I think Burge phrases a point more clearly elsewhere—I focus almost exclusively on *Origins* below, in no small part because this text incorporates many of the arguments from the earlier papers, but puts them in the context of the broader argument of *Origins*. Some points from these earlier papers are also explicitly developed further in *Origins* (cf. Burge, 2010, p. 545fn), and so supersede the earlier discussions.

³I use the term ‘constitutive’ and its cognates, as in the phrase ‘constitutive conditions’ almost exclusively, rather than using related terms such as ‘essentially’, or ‘fundamentally’ and their cognates, even where grammatically it would be less awkward to do so, because Burge gives these terms quite precise meanings, and I intend to follow his usage here to avoid confusion (Burge, 2010, pp. 57–59). To make the writing somewhat less repetitive I try to use the phrases ‘necessary condition(s)’ and ‘sufficient condition(s)’ and their cognates where appropriate, though in many cases this is not possible.

But this view faces an obvious problem. It, like all versions of individual representationalism, according to Burge, denies that infants and non-rational animals are capable of objectively representing, for although such creatures have sensations caused by mind-independent objects, they are incapable of thinking about such objects as the causes of their sensations.

One might suppose that intellectualism is *per se* a form of individual representationalism, and hence that intellectualist theories are *ipso facto* subject to this objection. However, in this chapter I argue that Burge's criticisms do not apply to intellectualist theories of perception that endorse what Matthew Boyle (2016) calls a *transformative theory of rationality*. According to transformative theories of rationality, "the very nature of perceptual and desiderative capacities [is] transformed by the presence of rationality, in a way that makes rational perceiving and rational desiring essentially different from their merely animal counterparts" (Boyle, 2016, pp. 530–531).

There are two central arguments in this paper, one modest and one ambitious. The most modest argument of the paper simply aims to show that Burge's arguments fail to gain traction against intellectualist views that endorse a transformative conception of rationality, for such views are not committed to denying that infants and non-rational animals are capable of objective perception.

In short, the argument is that the claims that intellectualists make about mature human perception imply that non-rational animals and infant humans are incapable of objectively perceiving only on the assumption that the necessary conditions for mature humans to be capable of objectively perceiving are also necessary for infants and non-rational animals to be capable of objectively perceiving. But this inference need not go through if we have a transformative conception of rationality. If the acquisition of rationality transforms those

capacities we share with non-rational animals, and with our infant selves, then there may be more, or more demanding, necessary conditions for rational creatures than for non-rational creatures to be in states that non-rational creatures may also be in. Being in an objective perceptual state may thus require more of a rational creature than a non-rational creature, simply because she has more sophisticated capacities than a non-rational creature does.

Although the conclusion of this argument is interesting in its own right, it is of principal relevance to the rest of the chapter because it clears the ground for the more ambitious argument. This argument builds on another aspect of Burge's view to provide an argument for a particular intellectualist view of perception. Specifically, Burge claims that a constitutive condition for objective perception is that objective perceptual states must be attributable to individuals, rather than merely to some of their subindividual systems. I argue that *rational* individuals have distinctively *rational* perspectives, and that this has the consequence that if we are dealing with a rational perceiver, the condition that perceptual states are attributable to individuals can only be met by perceptual states that are capable of standing in rational relations with her beliefs and other propositional attitudes, and of serving as reasons for belief.

Perceptual states that are not integrated with a perceiver's rational capacities to at least this degree may be states of subindividual systems, or of non-rational creatures, but they cannot be states of rational creatures. And so, given Burge's argument that objective perceptual states are constitutively attributable to individuals, such states would not be objective. For rational creatures, a constitutive condition for a perceptual state to be objective is that it be capable of standing in rational relations and serving as a reason for belief. Or so I will argue.

2.2 Tyler Burge on Objectivity

2.2.1 Second-Family Individual Representationalism

The topic of Burge's book is *objectivity*; specifically objective *representation*. Given that we will be concerned principally with perception, our focus will be on objective *perceptual* representation. Perceptual states, on Burge's view, are *sensory* states that objectively represent the physical environment and, hence, have *veridicality conditions* that depend on the features of the physical environment that are so represented.⁴ Thus, an objective perception of a red apple is a sensory state that represents that apple as red, and it has veridicality conditions that are determined by such a content. A perceptual state that represents an apple as red is veridical if the apple is red, and unveridical if either there is no apple causally connected to the state in the right way or the relevant apple is not red.⁵

Burge has his own view about objective representation, which we will consider in §2.2.2. However, a large part of *Origins of Objectivity* is devoted not to the exposition of Burge's view but rather to the attempt to refute alternative views. These views will be our concern in this section. His main target is a family of views about what makes objective representation possible. These views think of objective representation, whether perceptual or not, as a

⁴As Burge uses the terms, perception and representation are by definition objective (see §2.3.2 below). Thus, I will sometimes simply drop the term 'objective' altogether and speak of representation or perception *simpliciter*. Where there is the potential for confusion, such as when talking about the views of philosophers who use terms such as 'perception' and 'representation' to apply more broadly, I will try to be explicit, even if sometimes this leads to slightly more clunky prose.

⁵See, e.g., Burge, 2010, p.38ff. The same applies, *mutatis mutandis*, for capacities, acts, and events. I follow Burge in assuming that perception is representational, and that to be representational is to have veridicality conditions. The view I develop later in the dissertation is representational in a broad sense, in that it attributes content to perception, though the content so attributed is in the first instance articulated in terms of its inferential role rather than in terms of its truth- or accuracy-conditions. There is thus a narrower sense of 'representation' on which inferentialism is contrasted with representationalism about content, in which case the view does not count as a version of representationalism. For more on this, and on the connection between contents and accuracy or veridicality conditions, see §4.2.2.

fairly sophisticated achievement available only to rational, language-using creatures. Burge characterises his target more precisely as follows:

“According to this view, an individual cannot objectively and empirically represent a physical subject matter as having specific physical characteristics unless the individual has resources that can represent some central conditions under which such representation is possible. Individuals qualify as engaging in objective empirical representation by having resources for explaining what they are doing.” (Burge, 2009, p. 287)

Burge calls philosophers who accept this assumption *individual representationalists*. He gives a number of examples of capacities that individual representationalists have argued are necessary for a creature to be capable of objectively representing:⁶

“Examples of supplementary capacities are a capacity to represent a distinction between appearance and reality, or a capacity to represent laws or causal generalizations, or a capacity to represent criteria for identity or individuation.” (Burge, 2010, p. 16)

Burge acknowledges that the *existence* of a distinction between appearance and reality, of laws (or at least law-like regularities), and of criteria for identity or individuation are all constitutive conditions for objective representation. If there were no difference between how things seem and how things are, if there were no laws (or law-like regularities), or if objects had no criteria for individuation then perception would not be possible.⁷ However, individual representationalists go further than this. They argue that if an individual is to be capable of objectively representing it is not enough that these constitutive conditions

⁶See Burge, 2010, p. 19 for a more extensive list.

⁷This first claim, unlike the other two, may actually be controversial. Surely if we could simply see everything as it is, without the intervention of appearances, our representations would still be objective! As we will see below, Burge thinks that a distinction between appearance and reality is necessary for objective representation because he thinks that objective representation is only possible for creatures with sensory (principally visual) systems that can recover information about distal objects from impoverished stimuli. And he thinks of this process as, in effect, recovering properties of objects from appearances. (Though see Adams and Maher, 2016 for a challenge to this idea.) Nothing hangs on whether Burge is right to concede this point, however. Either way he denies the individual representationalist claim that a capacity to represent this distinction is necessary for objective representation.

are in fact satisfied. In addition, the individual must be capable of representing them; paradigmatically, she must be capable of *thinking about* them.

Burge distinguishes between what he calls *first* and *second family* versions of individual representationalism. The two families differ from each other both over the kinds of capacities they claim are required for objective representation and over the role that these capacities play in securing objective representation (Burge, 2010, pp. 13–22). First family views are rooted in the empiricist tradition, and tend to emphasise the idea that objective representation combines the deliverances of simple capacities for sensory responsiveness to the environment—e.g., ideas, sensations, or sense-data—with sophisticated capacities to conceptualise the relationship between these things and the environment—e.g. to think of one’s sensory states as being *caused by* objects, or as part of a pattern of sensory stimulation governed by sophisticated laws.⁸ Given that first family views are typically driven by an additive theory of rationality, we will not concern ourselves with them in what follows. I mention them here simply to bring out the pertinent features of second family views, on which we will focus from here on in.

Second family views are rooted more firmly in the rationalist tradition and so tend to reject the empiricist idea that objective representation depends on an isolable and more basic form of perceptual sensitivity to the environment.⁹ However, even if they do not think that rational capacities *supplement* a more basic form of sensory awareness, as first family individual representationalists do, second family views nonetheless insist that objective rep-

⁸Burge treats Russell as the paradigmatic first family individual representationalist (Russell, 1905, 1912/1972), citing specifically his view of knowledge by description (Burge, 2010, p.118ff), but Burge also mentions Moore, Broad, Price, Ayer, Schlick, William James, and early Carnap (Burge, 2010, p. 16).

⁹The association of empiricism and rationalism with first and second family views, respectively, is complicated by the presence of empiricists such as Quine in the second family. Furthermore, Burge suggests that individual representationalist assumptions crept into 20th century empiricism through the lingering influence of rationalism due to the influence of Frege (Burge, 2010, p. 116). However, the connection is still instructive, as long as these caveats are kept in mind.

resentation depends upon such capacities. Perception is not merely sensory responsiveness *plus* rational capacities, but neither is it autonomous with respect to such capacities (Burge, 2010, p. 16). Such views thus tend towards a kind of *holism*, arguing that the representation of ordinary macro-physical objects must be part of a broader network of capacities for more sophisticated forms of thought—e.g. a capacity for self-conscious thought, the ability to unify representations into a coherent view, or even a grasp of the concept of objectivity itself (cf. Burge, 2010, p. 19).¹⁰

Burge divides those he identifies as second family individual representationalists into two groups: neo-Kantians such as Strawson and Evans, on the one hand, and more linguistically-focused philosophers such as Davidson and Dummett on the other. It will be helpful to focus on one view here to illustrate Burge’s criticism. Although he is not one of Burge’s central targets, Sellars is perhaps the best figure to choose for this purpose, given that his view contains features of both neo-Kantian and linguistic versions of second family individual representationalism, and given that Sellars’s account is one inspiration for the intellectualist account developed later in this dissertation. Sellars acknowledges a sensory element in perception, but denies that this element is representational. Rather, as Gupta puts it, “the object-directedness of experience, according to Sellars, is at bottom the intentionality of thought” (Gupta, 2019, p. 37). Thus, Sellars says:

“coming to see something as red is the culmination of a complicated process which is the slow building up of a multi-dimensional pattern of linguistic responses (by verbal expressions to things, by verbal expressions to verbal expressions, by meta-linguistic expressions to object-language expressions, etc.) the fruition of which as conceptual occurs when all these dimensions come into play in such direct perceptions as that this physical object (not that one) over here (not over there) is (rather than was) red (not orange, yellow, etc.).” (Sellars, 1963, p. 334)

¹⁰Burge mentions Frege, Cassirer, Sellars, Dummett, Strawson, Evans, Quine, and Davidson as proponents of second family individual representationalism (Burge, 2010, p. 18). For some of the central arguments, see in particular Davidson, 1982; Evans, 1982; Frege, 1918; Sellars, 1963, 1956/1997; Strawson, 1959, 1966.

As the first parenthetical remark in this quotation indicates, Sellars thinks that the capacity to see (this) something (here) as red depends upon the capacity to use language, to respond to other users of language, and even to talk about language itself. Sellars thus suggests in this quotation that the ability to objectively perceive—to see something *as* one way rather than another—depends on such sophisticated linguistic abilities. So it looks like Sellars, in this passage, is denying that non-human animals and infants are capable of objective perception.¹¹

Burge accuses individual representationalists of *over-intellectualising* perception (Burge, 2010, pp. 257, 283, 285). And he claims that views, like Sellars’s, that claim that *language* is required for objective perception, are “perhaps the most hyper-intellectualized” (Burge, 2010, p. 283) of all of them. The philosophical substance of this accusation comes from Burge’s argument that the role such views give to intellectual capacities in perception conflicts with our best scientific accounts of perception:

“Most discussion of objective representation postulates conditions that are much more sophisticated and intellectual than are warranted. Nearly all second-family views leave it doubtful that animals and human infants perceptually represent (or represent as of) bodies, or any other element in the physical environment. Such positions... are overrun by perceptual psychology. Armchair arguments for Individual Representationalism are not nearly powerful enough to show... psychology to be mistaken.” (Burge, 2009, pp. 290–291)

The claim that intellectualist views of perception such as Sellars’s are, as Richard Rorty puts it, “unfair to babies” (Rorty, 1979, p. 181) is not a new one. (Rorty’s admittedly less subtle characterisation of this problem predates Burge’s by over 30 years.) What is unique

¹¹I choose this quotation in part because it is one of the few pieces of textual evidence that Burge actually gives in support of his claim that Sellars is an individual representationalist (see 2010, p. 139 for Burge’s brief discussion of this claim). The similarity with linguistic accounts is obvious, although the Kantian inspiration is perhaps not. Shortly after the passage just cited, however, Sellars says that the idea that seeing as depends upon initiation into a linguistic community, in the way described in this quotation, is “an adaptation rather than a rejection of Kant’s contention that the forms of experience are *a priori* and innate” (Sellars, 1963, p. 334). The Kantian credentials of Sellars’s view are expounded at length in *Science and Metaphysics* (1968).

about Burge's approach is not so much the accusation itself but the way Burge grounds it in a well-developed positive conception of perception based on a detailed understanding of perceptual psychology, as well as the way that he deploys this accusation against his opponents.

I will have more to say about Burge's positive view in the next section, but before we move on I want to say a little more about why Burge's deployment of this argument is novel. Although individual philosophers such as Sellars, and sometimes even whole traditions such as neo-Kantianism, have been accused of over-intellectualising perception, the novelty of Burge's approach comes from the breadth of the views that fall under the individual representationalist banner. Individual representationalism is not a view so much as an assumption shared by a number of views that are in many other respects astonishingly different from one another.

In effect, one of the aims of the argument in this chapter is to try to show that there are important differences between at least certain versions of intellectualism and other views that Burge classifies as versions of individual representationalism. If I am right then the crucial assumption is one that Burge shares with many of the individual representationalists he targets; namely, an *additive* view of rationality. If this assumption is rejected then Burge's accusations do not stick to many of his intellectualist opponents, even if they are accurate characterisations of some of his other targets.

2.2.2 Objectivity and Objectification

Arguably the most distinctive claim made by Burge in *Origins of Objectivity* is his claim that perception is a *natural kind*.¹² Burge argues for this claim by establishing three other claims:

“Three primary themes of the book are that objective representation is the basic sort of representation, that objectivity and representation begin in perception, and that perception is a very widespread and primitive capacity, present in numerous animals other than human beings.” (Burge, 2010, p. 10)

The first of these claims distinguishes Burge’s view from naturalistic approaches to perception that attempt to reduce objective representation to some other, more basic, kind of information transmission. Burge thus uses the term ‘representation’ in a stricter sense than many cognitive scientists and naturalistic philosophers of mind, who often apply the term to almost any information-carrying state. This is because not all states that carry information have veridicality conditions. Under normal conditions, smoke carries the information that there is fire, but it is implausible to suggest that the smoke (or the event of there being smoke) is veridical just in case there is fire.¹³

Burge argues that the sensory states of very simple animals no more represent their causes than smoke represents fire. Indeed, as we will see below in §2.3.2, it is crucial for Burge that we draw a sharp line between animals with very simple sensory systems and those capable of having genuine perceptual states with veridicality conditions. Only states

¹²As we will see below, Burge assumes that kinds can be specified in terms of necessary and sufficient conditions. This conception of kinds is wildly at odds with the orthodox view of biological and psychological kinds in the philosophy of science, where the Homeostatic Property Cluster (HPC) view is the default. See Taylor, forthcoming for a plausible attempt to draw out the consequences of the HPC view of kinds for Burge’s view of perception.

¹³Much of the history of late 20th century naturalistic philosophy of mind is devoted to the attempt to spell out what distinguishes mere natural meaning such as this from representation proper. For more on this see §2.3.2 below.

that *objectively* represent—that is, have accuracy conditions—represent at all, according to Burge. Merely reliably indicating, or carrying information, is not enough.

This brings us to the second claim made by Burge in the quotation above: The primacy of *perception* in the explanation of objectivity. According to Burge, the fundamental task of perceptual psychology is to explain how the visual system solves the *underdetermination problem*. The underdetermination problem arises because patterns of proximal stimulation received by perceiving creatures (often vastly) underdetermine their possible distal causes. This is almost trivially true in the case of vision, given that the retinal array is two-dimensional, whereas the visual environment is three-dimensional.¹⁴ The question that perceptual psychology aims to answer, therefore, is how the visual system manages to determine the actual causes of sensory stimulation, given the underdetermination problem. Only a creature whose visual system is capable of doing this can plausibly be said to be capable of objectively representing its environment.

A large portion of *Origins of Objectivity* is devoted to spelling out, in detail, the various ways that sensory systems (particularly the visual system) do just this. Given that I do not intend to dispute Burge’s characterisation of the science, we do not need to go into the details here (though see especially Burge, 2010, pp. 342–366, 396–416). The upshot, simply put, is that sensory systems solve the underdetermination problem by processing the information registered by sensory receptors in accordance with principles that allow the visual system to determine the likely causes of those patterns of sensory stimulation. Burge calls processes that solve the underdetermination problem in this way *objectifying* processes.

¹⁴Obviously the situation is more complicated than this, given that vision is typically binocular. The use of this fact is, however, one of the means by which the visual system solves the underdetermination problem. For more detail see Burge, 2010, pp. 347–350.

One conclusion Burge draws, then, is that a necessary condition for a sensory state to be representational is that it is the result of an objectifying process.¹⁵

Very simple organisms do not have objectifying sensory systems. Their responses can be explained entirely in terms of properties of the stimulation they receive, rather than in terms of representation of the causal antecedents of such stimulation:

“Bacteria can discriminate light from dark and move to where there is a paucity of oxygen. Nothing in their sensory capacities segments out the entities that have the attributes that they need (oxygen-poor water)... The organisms simply react to conditions on their bodily surfaces. These reactions are reliably correlated with environmental conditions that fulfill the organism’s needs. Negative response to light is regularly correlated with oxygen poverty... But bacteria do not perceive anything. They simply react to proximal stimulation. They sense light. They register it. They respond. Explanations of the capacities of bacteria are not enriched by invoking perceptual reference or perceptual attribution.” (Burge, 2010, p. 325)

By contrast, sensory systems that operate according to objectifying processes distinguish, in effect, between those aspects of sensory stimulation that are relevant for determining the properties and arrangement of physical objects in the environment and those that are not. One can only explain the proper functioning of these processes in terms that appeal to the representation of things in the individual’s environment. As Burge puts it:

“Perceiving involves having certain subindividual competencies that systematically filter contextually idiosyncratic elements in a signal from elements likely to bear on environmental reality.” (Burge, 2010, p. 282)

We can put this somewhat anthropomorphically by saying that objectifying processes distinguish between subjective and objective in the way that they process the information they receive from sensory receptors. Objectifying perceptual processes are thus deserving

¹⁵I say *a* necessary condition because, as we will see below, Burge also thinks that representational states must be attributable to the individual. And I say *sensory* state because a non-sensory state might be representational not because it is the result of an objectifying process but because it is related in the right kind of way to a sensory state that is (see, e.g., Burge, 2010, p. 3).

of the name, and indeed are the primary source of objective representational states more generally. Whilst there are objective representational states that are not directly generated by objectifying perceptual processes, it is only because there are such processes that any states have representational content at all. Non-sensory objective states such as beliefs get their status as objective indirectly, by means of their integration with representational states that are the result of objectifying processes (Burge, 2010, pp. 3–4).

This point, finally, brings us to the last of Burge’s three claims made in the quotation above: His claim that perception, in the full blooded sense just described, is a widespread and primitive capacity. This claim distinguishes Burge’s view from individual representationalist approaches, which *deny* that objective perception is widespread, for reasons that we saw above.

Although objectifying processes are quite complex, the principles that govern their operation are, as the quotation above makes clear, *subindividual*. Thus, the creature itself need not be capable of understanding the objectifying principles that make perception possible for them. Burge even insists that the visual system itself need not be sophisticated enough to have states that have *as their content* the principles that govern the proper functioning of the system. It is enough that these principles *in fact* govern the processing of sensory informational states:

“Objectification is a *subindividual* process. Moreover, the separation of proximal stimulation from purportedly system-independent items in the environment is *represented* neither by the individual nor by the system. In perception, there is no representation of any distinction between subjective and objective.” (Burge, 2010, p. 402)

Thus, even quite cognitively simple creatures may have objectifying visual systems, and hence may be capable of objectively representing. It is this conception of perception and objectivity that underlies Burge’s accusation that individual representationalists hyper-

intellectualise perception, and which thus distinguishes Burge’s view from more simplistic versions of this criticism.

Thus, for these reasons, Burge argues, individual representationalism should be rejected because it conflicts with our best scientific understanding of perception (Burge, 2010, p. 275). Representation is more sophisticated, and less widespread, than many vision scientists and reductionist naturalistic philosophers of mind have supposed, for genuine representation involves more than just information transmission, but nonetheless it is still much *less* sophisticated, and *more* widespread, than individual representationalists suppose, for it does not depend in any way on sophisticated capacities only had by rational creatures.

2.2.3 Transformative Conceptions of Rationality

In the next section I want to suggest that even if many intellectualists have fallen into the trap of denying that infants and non-human animals are capable of objectively representing, a more plausible version of intellectualism about perception *does not* deny perception to non-rational animals, even though it characterises the perception of *rational* animals in distinctively rational terms. In order to develop this idea, however, we first need to explain the difference between transformative and additive theories of rationality, which is the task of this section.

The kind of approach I am proposing is characterised by Matthew Boyle as a ***transformative theory of rationality***. Transformative theories of rationality give a distinctive characterisation of the connection between rationality and perception, which Boyle characterises as follows:

“our rationality... gives us a ‘special form’ of perceptual sensitivity to our environment, one whose operations are themselves informed by our capacity to weigh reasons... [What] ‘perception’ signifies in the case of rational creatures cannot be explained without reference

to the capacity for rational reflection. . . [Therefore,] an account of our sort of perceiving must itself appeal to capacities connected with rational thought and judgment.” (Boyle, 2016, p. 530)

Boyle’s characterisation of rational perception ought to strike us as remarkably similar to the characterisations given above of individual representationalism. And, indeed, elsewhere Boyle gives as examples of “capacities connected with rational thought and judgment” the capacity to think about oneself in the first person, which is one of the capacities mentioned explicitly by Burge as a characteristic ‘supplementary’ capacity proposed by individual representationalists as a necessary condition on the possibility of objective representation.¹⁶

Transformative theories of rationality are thus intellectualist theories, in the sense defined above. But Boyle is also insistent that transformative theories of rationality *do not* deny that non-rational animals can perceive, or that they can represent. Rather, transformative theories claim that we should understand rational perception and non-rational perception as two species of the genus *animal perception* (Boyle, 2016, p. 531). As John McDowell puts it: “we have what mere animals have, perceptual sensitivity to features of our environment, but we have it in a special form” (McDowell, 1996, p. 64).¹⁷ So transformative theories are not forms of individual representationalism. They do not deny that non-rational animals perceive; they simply deny that non-rational animals perceive *in the same sense that we do*.

It is thus important to distinguish between individual representationalist and transformative versions of intellectualism.¹⁸ Intellectualists all agree that we cannot explain the

¹⁶See Burge, 2010, p.157ff for Burge’s discussion of Strawson’s appeal to self-consciousness as a necessary condition of this kind. For Boyle’s discussion of first-person thought and its connection to transformative conceptions of rationality, see especially ‘Kant and the Significance of Self-Consciousness’ (2005).

¹⁷For Boyle’s discussion of this quotation, see Boyle, 2016, p. 531.

¹⁸Boyle speaks fairly exclusively of *theories* of rationality. Given that Burge doesn’t have an explicit *theory* of rationality, I will sometimes also speak more generally of transformative *conceptions* of and *approaches* to rationality. I will also speak of transformative theories or views of *perception*, where this means a theory or view of perception that presupposes a transformative conception of rationality. The same applies, *mutatis mutandis*, to additive theories.

perception *of rational creatures* without reference to the capacity for rational reflection. If one thinks that such capacities are required for objective perception *tout court* then one will, *a fortiori*, think that they are required for the perception of rational creatures in particular. So second family individual representationalist theories are versions of intellectualism.¹⁹ But not all intellectualist theories are versions of individual representationalism. If one holds a transformative intellectualist theory of perception one will only take rational capacities to play a constitutive role in the perception of *rational* creatures.

Below I will discuss the implications of this distinction for Burge’s criticisms of intellectualist approaches to perception. But first we need a characterisation of the general features of *non*-transformative conceptions of rationality. Boyle contrasts transformative theories of rationality with ***additive theories of rationality***. On an additive theory of rationality,

“rational animals perceive and desire in the same sense in which any animal perceives and desires; the power that differentiates our minds is something separate and additional... [Thus, according to additive theories,] an account of our minds might begin with an account of what it is to perceive and desire, in terms that do not presuppose the capacity to reflect on reasons, and then supplement this with an account of the ‘monitoring’ and ‘regulating’ of belief-on-the-basis-of-perception and action-on-the-basis-of-desire that only rational creatures can perform.” (Boyle, 2016, p. 528)

Thus in contrast to transformative theories of rationality, additive theories claim that the ability to reflect on reasons and to base beliefs on perception is an *additional* capacity, not necessary for perception itself, but only for more sophisticated abilities unique to rational creatures. According to additive theories, therefore, we can provide an explanation of the capacities we share with non-rational animals, such as perception, *without* reference to the capacity for rational reflection.

¹⁹First family individual representationalist views are not versions of intellectualism because they posit different supplementary capacities. But such views are also committed to an additive conception of rationality, conceiving as they do of these capacities as supplementing a more basic form of perceptual sensitivity. So the defense below does not apply to them.

Note that transformative theories need not deny that there is *some* common characterisation that can be given of rational and non-rational animal perception. Clearly there are likely to be physiological similarities between the visual systems of rational creatures and those of non-rational creatures. Transformative theorists can thus agree with additive theorists that it may be possible to give an account of these common features without having to make reference to capacities unique to rational creatures. The claim that transformative theories make, and that additive theories *deny*, is just that an account of these common features is not yet an account of the *perception* of rational creatures.

Given that the point is about whether or not we have to appeal to rational capacities in our *account* of perception, we thus need to note an ambiguity in the way Burge characterises individual representationalist theories. Burge often characterises individual representationalism as the claim that to be capable of objectively representing a creature must *have* capacities to represent some of the constitutive conditions for objective representation. Putting the point in terms of mere possession of capacities it is difficult even to articulate the claim made by many transformative intellectualist approaches in a non-trivial manner.

The closest we can get in these terms is to say that for a rational creature to perceive it must *have* rational capacities, even though this is not a necessary condition on the perception of non-rational creatures. But this is a trivial claim, for rational creatures, in the sense we are using the term, *just are* creatures with rational capacities. Even Burge would be happy to accept that it is a necessary condition for a *rational* creature to perceive that it have this capacity, but this is just because having such capacities is necessary for being rational in the relevant sense.

At other times, however, Burge characterises individual representationalism in a more specific way. He says, for example, that individual representationalists claim that percep-

tion “constitutively depends on” (Burge, 2010, p. 257) capacities to represent some of the conditions for objective representation, or that representations “incorporate within themselves” (Burge, 2010, p. 13), or are “infused with” (Burge, 2010, p. 433) rational capacities. These and similar locutions highlight that some individual representationalists—most intellectualists—claim not merely that a creature must *have* rational capacities, but that such capacities are, in some sense to be specified more precisely, *involved* in perception itself. The fact that Burge takes these two formulations as roughly equivalent is, I think, an indication of the fact that he simply assumes that all of the views he targets, including intellectualist theories, hold an additive conception of rationality.

Borrowing terms from Boyle’s characterisation of transformative theories, we can put the point by saying that according to transformative versions of intellectualism we cannot *explain* what ‘perception’ signifies as this term applies to rational creatures without reference to rational capacities, even if this is not true of non-rational animals. Below, in §2.4, I aim to cash out these metaphors in more concrete terms in providing a positive transformative view of perception, and showing what role rational capacities play in the explanation of objective perception in rational creatures. But first, I put this distinction to work in the next section to address Burge’s criticisms of intellectualist views of perception.

2.2.4 Burge’s Blind-Spot

Burge does not explicitly endorse a theory of rationality as such. However, in this section I want to argue that many, if not most, of Burge’s arguments against intellectualist theories of perception presuppose, rather than argue for, an additive theory of rationality. If this is correct then his arguments beg the question against transformative versions of intellectualism about perception. Transformative versions of intellectualism do not conflict with

vision science because the apparent conflict between intellectualism and perceptual psychology depends on the assumption that we should think of rationality additively, rather than transformatively, and this is a philosophical assumption made by Burge, not a consequence of the relevant science.

The central indication that Burge assumes an additive theory of rationality is that he frequently draws conclusions about the conditions sufficient for rational animals to perceive from discussions of conditions that are sufficient for non-rational animals to perceive. But only by assuming an additive theory of rationality can we draw immediate conclusions about the sufficient constitutive conditions for *rational* animals to perceive from observations about the sufficient constitutive conditions for *non-rational* animals or infant humans to perceive.²⁰ If a transformative theory of rationality is correct then conditions that are sufficient for non-rational animals to perceive may be *necessary* for rational animals to perceive but they will not be sufficient.

For example, in a representative passage discussing the connection between objective representation and propositional capacities, Burge says the following:²¹

“[The] capacity to represent bodies as such [viz., objective representation]... phylogenetically and developmentally precedes, and *hence* is constitutively independent of, language and thought.” (Burge, 2010, p.437, my emphasis)

Burge evidently thinks that it follows from the claim that non-human animals and infant humans are capable of engaging in objective representation that objective representation is *in adult humans* constitutively independent of capacities unique to adult humans such as language or thought. And many of Burge’s other criticisms of intellectualist accounts of

²⁰I am assuming that the distinction between additive and transformative theories is exhaustive. What matters for the argument that follows is simply that Burge assumes that rationality is *not* transformative, so it ought not affect the argument if there is a third, unexplored view that lies between the two.

²¹In addition to this quotation, see Burge, 2010, pp. 195, 205, 209, 343 for a fairly representative sample of similar arguments.

perception take this same form. But if the acquisition of rationality transforms those capacities that we share with non-human animals (and our infant selves) then these observations do not undermine the claim that objective representation *in adult humans* is constitutively dependent on language, thought, or other rational capacities.

In other words, one consequence of adopting a transformative theory of rationality is that one cannot specify sufficient constitutive conditions for perception *tout court*. If the transformative theory of rationality is correct then the sufficient conditions for non-rational animals to objectively perceive may be *necessary* conditions for rational creatures to objectively perceive but they will not be *sufficient*, because our explanation of the perception of rational creatures will make essential reference to rational capacities not shared with such creatures.

Burge is right, however, to criticise intellectualists for underestimating the sophistication of non-human animal and infant perception. Many of the intellectualists Burge criticises assimilate the perception of non-human animals to mere sensory stimulation, or endorse some kind of crude behaviourism about animal perception.²² This assumption unfortunately colours intellectualists' response to concerns about the implications of intellectualism for non-rational animal perception. Take, for example, the following characterisation that Rorty gives of Sellars's response to the 'unfair to babies' objection I mentioned earlier:

“The existence of raw feels—pains, whatever feelings babies have when looking at colored objects, etc.—is the obvious objection to [Sellars's view]... To counter this objection, Sellars invokes the distinction between awareness-as-discriminative-behavior and awareness as what Sellars calls being ‘in the logical space of reasons, of justifying and being able to justify what one says’ [Sellars, 1956/1997, §36]. Awareness in the first sense is manifested by rats and amoebas and computers; it is simply reliable signaling. Awareness in the second sense

²²I focus on Sellars below, but similar examples are not hard to find in other figures targeted by Burge. Davidson, for example, claims that without triangulation between speakers “there would be no grounds for selecting one cause rather than another as the content-fixing cause”, and hence that non-linguistic animals do not have contentful states in this sense (Davidson, 1990, p. 203).

is manifested only by beings whose behavior we construe as the utterance of sentences with the intention of justifying the utterance of other sentences.” (Rorty, 1979, p. 182)

As formulated, Sellars (and, by implication, Rorty) seemingly falls into the trap of individual representationalism as a result of his intellectualistic view of adult human perception. Rorty characterises the awareness of rats in terms of merely reliable signalling, akin to the kind of awareness mere amoebas have, contrasting this with the linguistically- and conceptually-informed awareness of adult humans. It thus looks as if Burge is correct to accuse Sellars of “[moving] quickly from a non-representational notion of sensing to a propositional sensing-that... with no room for any type of perception in between” (Burge, 2010, p. 139fn).

Insofar as this pessimism about non-human animal and infant perception is an essential feature of intellectualism, I am happy to concede with Burge that it should be rejected, along with other forms of individual representationalism. But if I am right, however, then the intellectualistic part of such views is separable from the pessimistic view of the capacities of non-rational creatures with which it often comes packaged, and on the basis of which it is sometimes motivated. What matters, for example, according to Rorty’s characterisation of Sellars’s view above is *not* the overly simplistic characterisation he gives of non-rational awareness, but simply the contrast he claims Sellars draws between awareness of this kind and the kind of awareness had by rational animals.

The contrast between these two sorts of awareness is no doubt particularly marked if, as Rorty suggests, non-rational creatures merely reliably indicate rather than objectively represent. But there may still be an important distinction to be drawn between the kind of awareness had by non-rational creatures and the kind of full-blooded rational form of

awareness that is Sellars's focus, *even if Burge is right to insist that non-rational animals are capable of objective perception in a philosophically substantive sense.*

What matters, for this kind of defence to succeed, is just that the awareness had by rational creatures differs in kind from that of non-rational creatures. And this can be true even if Burge is right to insist that the kind of awareness had by non-rational creatures is genuinely objective perception, and not mere reliable signalling. The only point intellectualists are committed to is that the constitutive conditions for objective perception may differ for rational and non-rational creatures. And if a transformative theory of rationality is true, we have an obvious explanation for why this should be so: Both rational and non-rational creatures are capable of objectively perceiving, but the explanation of objective perception *in rational creatures* necessarily involves appeal to their rational capacities.

In other words, intellectualists who *do* fall into the trap of individual representationalism may be right to insist that there is a difference in kind between the perception of rational and non-rational creatures, even if they are mistaken in characterising this as a difference between objective and non-objective perception. This, of course, does not yet give us a reason to suppose that a transformative view of rationality is correct, but it does show that *if* we endorse a transformative theory intellectualism need not have the implications Burge claims such views do.

It is telling, in connection with this point, that all of Burge's main targets are philosophers writing predominantly in the middle of the 20th century or earlier, before much of the science cited by Burge in support of his view was settled. Furthermore, Burge himself criticises vision scientists for conflating objective representation with mere information transmission, or 'reliable signalling', in Rorty's phrase. Thus, one can hardly fault Sellars, for example, for failing to credit non-linguistic animals with abilities that he had no evidence to believe

they had, even according to the consensus of the best scientists at the time he was writing.²³ The interesting issue is thus not whether these philosophers underestimated the capacities of non-rational animals but whether the contrast they draw between rational and non-rational animals rests *essentially* on this underestimation. If the argument I have given above is correct, it need not.²⁴

In this section I have only aimed to establish a relatively modest conclusion; namely, that intellectualism about perception clashes with vision science only if we assume an additive theory of rationality. On this basis, we can draw conclusions about the sufficient conditions for rational creatures to be capable of perceiving from sufficient conditions for *non*-rational creatures to be capable of perceiving. But if a transformative conception of rationality is correct then this inference is blocked. At least some of the necessary conditions for rational creatures to perceive will involve reference to those creatures' rational capacities. And this distinction may still hold even if the perception of non-rational creatures is still genuinely objective, as Burge argues it is.

This argument is modest for the obvious reason that I have not attempted to give any reason to think of rationality as transformative, or that the constitutive conditions for rational creatures to perceive make reference to such creatures' rational capacities. What we need, in other words, is an argument for a transformative intellectual account of perception. In the next section, I explain two aspects of Burge's positive view that, I argue, provide the foundation for a more ambitious argument of just this kind.

²³To his credit, Burge explicitly acknowledges this (Burge, 2003, p. 527). However, he does not take the extra step of considering whether intellectualists who *are* aware of contemporary science might be able to develop a view compatible with his view of objective perception. Since one can hardly fault Burge for this, either, part of the point of this chapter is to rectify this oversight.

²⁴At least some contemporary philosophers working within the tradition of Burge's targets explicitly endorse a transformative view of rationality, though Burge does not engage in detail with any of them in *Origins of Objectivity*. Burge does engage in a more detailed discussion with McDowell in a series of papers, however. See Burge, 2005, 2014 and McDowell, 2011a, 2011b, 2013b.

2.3 Perception and Explanation

2.3.1 Burge's Two Arguments

I noted above, in §2.2.2, that Burge criticises vision scientists for relying on a deflated conception of representation as mere sensory discrimination. One related aspect of Burge's view not discussed so far is his insistence that perceptual states are attributable to *individuals*, and not merely to some of their subsystems. Vision scientists, partly as a result of their deflated conception of representation, often take the primary bearers of content to be subindividual sensory systems, rather than individuals.²⁵ Even worse, perceptual psychologists often fail to pay close attention to this distinction altogether, moving indifferently between talk of visual systems manipulating representations and attribution of representational states to the creature itself. Burge insists, against both of these tendencies, that objective representational states are constitutively states of individuals.

Burge's story is thus by necessity somewhat subtle. As we saw in §2.2.2, he argues that a necessary condition for a creature to be capable of perceptually representing is that the creature have a sensory system capable of objectification. Burge also insists that objectifying processes are constitutively *subindividual*; this point forms part of the basis of his criticism of individual representationalists, who argue in effect that objective representation requires understanding of some of the principles governing such systems. Nonetheless, the perceptual states *to which these subindividual processes give rise* are constitutively attributable to *individuals* and not merely to the subsystems that give rise to them.²⁶ On Burge's picture,

²⁵I follow Burge in saying 'subindividual' rather than 'subpersonal' insofar as I am discussing the distinction as it applies both to rational creatures and non-rational creatures, given that the latter are not persons. In my discussion of rational individuals below, in §2.4.1, I revert back to the more familiar terminology.

²⁶Burge claims that they are *also* states of the visual system (Burge, 2010, p. 369)—what is important is simply that they are not *only* states of the visual system.

then, there are subindividual necessary conditions for something to count as a perceptual state, but only states attributable to individuals actually have representational contents.

Burge gives two arguments for the claim that perceptual states are constitutively attributable to individuals. The first concerns the relationship between perception and action. Burge argues that perceptual states are partly dependent for their representational content on interactions between the whole animal and its environment. In navigating, eating, attacking, and mating the creature fulfils various individual needs, and contributes to various whole-animal functions such as survival.²⁷ The environmental entities and properties with which the animal interacts in fulfilling such functions constitutively constrain the entities and properties representable by the creature (Burge, 2010, pp. 319–321). It is because the creature needs to eat *plants*, navigate around *rocks*, and attack other *animals* that it is objects of this kind, rather than galaxies or cells, that are represented in perception.

Conversely, perception constitutively figures in the explanation of individual activities and of the realisation of whole-animal function. Any satisfactory explanation of the behaviour of a cat as it navigates around obstacles in its environment in search of its prey must make appeal to the cat’s perceptual representation of macro-level objects and their properties. Given the constitutive connections between perception and individual activities, functions, and needs Burge argues that perception *must* be attributable to the individual creature itself. Were representational states merely attributable to some of the creature’s subsystems, it could not play this explanatory role. Perception is thus “constitutively a way of representing goals, obstacles and threats *for individuals*” (Burge, 2010, p.371, my emphasis).²⁸

²⁷I will follow Burge—e.g. Burge, 2010, p. 371—in using ‘whole animal’ and ‘individual’ relatively interchangeably. I will use the former in places where the singular connotations of ‘individual’ are unwelcome, such as in discussions of species-specific function.

²⁸The conjunction of these two arguments might initially seem circular: How can action *both* play a role in constraining the contents of perception *and* perception play a role in explaining such action? Burge dispels

Burge's second argument concerns the relationship between perception and the notion of a representational perspective. This argument is short, so we can quote it in full:

“The objectification and representational content involved in perception constitutes a point of view on environmental *representata*. Representational perspectives or points of view are constitutively attributable to individuals. Psychology tries to explain how the representational perspectives of individuals are formed. Perceptual kinds ground this type of explanation. To play this role, they must be attributable to individuals.” (Burge, 2010, p. 371)

It is difficult to know exactly what Burge means by ‘representational perspective’ or ‘point of view’. He says that the perspective of an individual is constituted by modes of representation, and that the latter “mark how the world is, representationally, for an individual” (Burge, 2010, p. 37). He also says that an individual creature is “a locus of perceptual representation” (Burge, 2010, p. 371). I will say more about this idea as I think it applies to rational creatures in §2.4.1 below, but for now this characterisation will suffice.

To summarise, Burge has two arguments to the effect that perceptual states are constitutively attributable to individuals. The two arguments are structured similarly. Burge argues that some features—actions in the first case, and perspectives in the second—are constitutively attributable to individual creatures. He then notes that perception plays a constitutive role in the explanation of this feature. He concludes that perception, if it is to play this explanatory role, must also be attributable to individual creatures, and not merely their subsystems.

this appearance of circularity by distinguishing primitive agency from other forms of activity. Primitive agency is possible even for creatures incapable of objectively representing. The whole-animal functions and needs of such creatures constrain the representational contents of such creatures and of creatures phylogenetically downstream from them. But once a species is capable of objective representation the kinds of activities in which creatures of that species are capable of engaging become more sophisticated. As Burge puts it: “When perception sets an object for animal action, agency reaches a new level of sophistication. The animal itself perceptually represents the goal that the action fulfills.” (Burge, 2010, p. 337) In *such* creatures, explanation of whole-animal function, needs, and activities will make appeal to representational content. See Burge, 2010, pp. 326–341 for an extended discussion of primitive agency. I think this is another instance of what I call below Burge's *transformative conception of objectivity*.

We have seen the evidence that Burge gives for the first premises of these two arguments—the claims that the relevant features are constitutively attributable to individuals. But what about the assumption that perception can only play a constitutive role in the explanation of an individual-level feature if it is attributable to the individual? Why couldn't *subindividual* states play a role in the explanation of features or behaviours of individuals?

It does not seem necessary, in order to argue that perceptual states are constitutively states of individuals, to deny that states not attributable to an individual might play a role in the explanation of some of its individual activities or features. Amongst other things, on Burge's own account subindividual perceptual processes clearly play an important role in the explanation of a creature's perceptual states which, according to Burge, are attributable to the individual.

The issue seems to be that subindividual states are not capable of featuring in the *right kind* of explanations; namely *psychological* explanations. A psychological explanation of a creature's behaviour will appeal essentially to states attributable *to the creature* to explain why it behaved in the way it did: We explain why the cat pounced by attributing *to the cat* perceptual awareness of the mouse in front of it.²⁹ We could also explain its behaviour in terms of the contraction of muscle fibres or the firing of neurons, all of which are subindividual, or even by appeal to biological function, but to do so would not be to give a *psychological* explanation. I explore this idea further in the next section.

Being a state that is the upshot of an objectifying perceptual process is thus a *necessary*, but not sufficient, condition for being a perceptual state. Such states must also be attributable to the individual creature in order for them to play the role they do in (primar-

²⁹We should not understand 'awareness' here as necessarily involving *consciousness* (though in conscious creatures it ordinarily will do), for Burge denies that states attributable to the individual need be conscious (Burge, 2010, pp. 374–376). Despite the brevity of Burge's discussion of this point (especially in comparison with the space devoted to other issues), this has proven to be one of the most controversial claims in the book. For what is, I think, a decisive treatment of the issue, see Taylor, forthcoming.

ily psychological) explanations of the individual's behaviour and perspective. In the next section I argue that Burge's view of *irreducibility* of the relevant notions of psychological explanation and function leads him to endorse a *transformative conception of objectivity*.

2.3.2 Burge's Transformative Conception of Objectivity

Above I argued that Burge presupposes an additive theory of rationality in his arguments against intellectualist theories of perception. Thus, Burge himself is unlikely to be sympathetic to views that endorse a transformative conception of rationality. In this section I want to argue that despite this fact an intellectualist sympathetic to Burge's account of objective perception might nonetheless find some leverage against Burge in the details of his own view, for Burge himself has a transformative conception of *objectivity*.

In the previous section I mentioned that Burge criticises vision scientists for using the term 'representation' in a deflationary way, to apply to any state of a creature that carries information about its environment. Burge criticises them not only for being sloppy about the distinction between individual and subindividual states as a result, but also because he claims that this elides important explanatory differences between states that merely carry information and those that actually represent. What is telling, however, is the way that Burge makes this point:

“representation and perception... are to be distinguished from other sorts of functional information registration—and, in the case of perception, other sorts of functional sensory information registration... ‘Representation’ is often used, in science and philosophy, to apply to such systems. I argue, on scientific grounds, for a narrower application. The point is to show that *representational mind is to be distinguished from other functional information systems. It constitutes a distinctive ‘species’ or kind—a ‘cut’ in nature.* Perception is situated just above the lower border of that ‘cut’. As noted, this border—which demarcates origins not only of perception, but also of representation and objectivity—begins at more

primitive levels than philosophy has traditionally recognized. ” (Burge, 2010, pp. xi–xii, my emphasis)

Importantly, Burge goes to significant lengths to reject an *additive* conception of objectivity, in terms of sensory discrimination *plus* some other condition. For example, some teleological accounts add to a deflationary account of sensory discrimination a *biological* notion of proper function, thereby hoping to explain the notion of objective representation and representational function in terms of biological function. Those who endorse teleological views of this kind typically try to account for (objective) representational accuracy in terms of a state’s fulfilling its proper function and to account for representational error in terms of its *failing* to fulfil its proper function.³⁰ Burge argues that such additive approaches to objectivity are misguided insofar as they see genuine representation as simply lying at one end of a continuum whose other end terminates with mere reliable sensory discrimination, rather than seeing the development of objective representational capacities as carving out a new joint in nature:

“There is nothing wrong with a broad notion of differential sensitivity and responsiveness that is associated with function. Such a notion describes the sensory capacities of many organisms, including the sensitivities of plants. One can use any term to express such a notion that one likes, including ‘representation’. Many biologists and psychologists, and some philosophers, do use the term ‘representation’ in this way. What I doubt about this tradition is an additional claim. This is the claim that the foregoing notion is the only scientifically respectable notion of representation *and that there is no significant difference in kind between this notion and any notion of representation employed in psychological explanations.*” (Burge, 2010, p.293, my emphasis)

³⁰Millikan (1984, 1989) is perhaps the central figure. Millikan’s account, like that of Papineau (1984), *replaces* the informational account with a consumer-based teleological one, rather than supplementing it, though it does still rely on a notion of sensory discrimination. Dretske’s (1986) account, by contrast, combines his earlier informational account (1981) with a teleological condition, though Dretske eventually settles on a conception of function determined by learning rather than biology. Obviously this summary is oversimplified, but given that the account is mentioned here simply as a foil for Burge’s view, I pass over the important differences.

In essence, that is, *Burge has a transformative conception of objective representation*. On Burge's account, the acquisition of capacities for objectification, and the distinctive conception of representational function and explanation that goes along with such capacities, is not simply an extra feature of the mind that we add to capacities for differential responsiveness. Rather, it fundamentally transforms the way a creature responds sensorily to its environment, constituting a new *kind* of awareness, with the implication that the term 'representation', and by extension 'perception', actually *means something different* when it applies to creatures who do have such capacities than it does when applied to creatures whose states merely carry information about their environment.

Crucially, for the purposes of the argument I consider below, Burge does not attempt to motivate his claim about a difference in kind on physiological grounds. Physiologically and anatomically "there is a continuum between an amoeba's sensitivity to light and human vision" (Burge, 2010, p. 319). The distinction in kind posited by Burge is motivated rather by the kinds of *explanations* in which these states play an essential role; namely, psychological explanations of the kind we discussed in the previous section.

According to Burge, genuinely representational states are states that feature in explanations of individual perspectives and behaviour that appeal *essentially* or *non-trivially* to the veridicality conditions or content of those states. Thus, whilst one can *describe* a thermometer as representing the temperature, this does not provide evidence that thermometers genuinely represent anything, for there is no *need* to appeal to veridicality conditions to explain the behaviour of a thermometer (Burge, 2010, p.292ff). We can give a perfectly sufficient explanation of the thermometer purely in terms of reliable covariance between the states of the thermometer and the temperature. The appeal to content in this case is an idle

wheel. By contrast, Burge thinks that the explanations of empirical psychology and vision science *essentially* make appeal to states with veridicality conditions.

Thus, by attempting to reduce the notion of representational function to biological function, teleological accounts conflate genuine representational accuracy with mere practical utility (Burge, 2010, p. 301ff). Psychological explanations are not reducible to explanations in terms of biological function, for explanations in terms of biological function make no *essential* appeal to veridicality conditions (Burge, 2010, p. 302). States with veridicality conditions are governed by a notion of *representational function* that cannot be reduced to the notion of proper biological function, for though perceptual accuracy typically contributes to fitness and survival, it *need not*. There is no constitutive connection between fulfilment of biological function and fulfilment of representational function (Burge, 2010, p. 302).

For this reason we cannot simply add *non*-representational conditions, such as a biological conception of proper function, to a *non*-representational conception of sensory discrimination and expect to arrive at an adequate conception of representation. To borrow a metaphor from Brandom (1994a, p. 41), Burge argues that you can't bake a representational cake out of non-representational ingredients.

I think Burge's arguments against these views are good, as is his transformative conception of objectivity. But Burge thinks of his account as providing not just a characterisation of the constitutive conditions for perception in non-rational animals, but as providing *the only plausible account of perception* for rational and non-rational animals alike (Burge, 2010, p. 295). It is this assumption that I want to challenge below.

Echoing the transformative language that Burge himself uses in his criticism of his opponents, we might say that whilst we can certainly use the term 'perception' in the way that Burge does, it is not the only respectable notion of perception. Burge shows that although

any creature with a sensory system is capable of sensory discrimination, we should not think of objective perception as perception, conceived as mere sensory discrimination, *plus* some extra condition that makes it objective. Rather, the explanations in which objective perceptual states figure make *essential* appeal to features not shared by sensory states that *merely* carry information about the environment.

Creatures not capable of objectively representing still perceive in a sense—they have sensory awareness of their environment—but creatures capable of objectively representing perceive *in a different sense*—we cannot explain their form of sensory awareness without appeal to the notion of veridicality conditions and content. Similarly, I want to argue, although any creature with an objectifying visual system is capable of objective perception, we should not think of rational perception as *mere* objective perception *plus* some extra condition. Rather, the explanations in which rational perceptual states figure make essential appeal to features not shared by objective perceptual states that *merely* objectively represent in a non-rational way.

2.4 Perception and Rationality

2.4.1 Individuals, Persons, and Reasons

According to Burge, perceptual states are constitutively attributable to rational creatures and constitutively play a role in the explanation of their perspectives and actions, but perception is nonetheless constitutively *independent* of those states and capacities unique to rational creatures, such as their propositional attitudes, their capacity to reflect on reasons for their beliefs, and their capacity to act on the basis of reasons (cf. Burge, 2010, p. 437,

quoted above in §2.2.4). The sufficient conditions for rational creatures to objectively perceive are the same as those for non-rational animals: Their perceptual states need to be the upshot of objectifying perceptual capacities, and they need to be attributable to perceivers as individuals. In this section I will argue that this combination of views is untenable.

The reason for this, I will argue, is that rational individuals are, in a sense I will spell out in more detail below, essentially individuals that bear propositional attitudes such as belief that are governed by norms of consistency and coherence: They are *thinkers*. For this reason, perceptual states can be attributed to such individuals, rather than merely to their subpersonal visual systems, only if these states stand in rational relations to their propositional attitudes, and can serve as reasons for belief.³¹ A perceptual state that was not capable of standing in such relations and serving as a reason for belief in this way would not be a state of such an individual. However, since we are not presupposing an additive theory of rationality, it might nonetheless still be the state of a non-rational individual, for these same constitutive conditions do not apply to the perspectives of non-rational individuals. In a slogan: *Perception* is transformed by the acquisition of rationality because the *perceiver* is transformed by the acquisition of rationality.

To see why this is so we should begin by considering in more detail what the relevant notion of a perspective is, as it applies to rational individuals. We saw above that Burge characterises representational perspectives in terms of the notion of a *point of view* on the world, and characterises an individual as a *locus* of representations. Burge does not say what he means by this, but here I want to provide an interpretation of these ideas as they apply to rational individuals.

³¹Although below I focus almost exclusively on belief, I think similar arguments apply also to other propositional attitudes. I do not try to give these arguments here because the form of the argument would be essentially the same, and so would simply take up space without adding anything of philosophical substance.

Whatever else they might be, rational creatures are bearers of distinctively *conceptual* representational states; principally propositional attitudes such as beliefs and judgements.³² A rational individual has, by virtue of being the bearer of such states, a distinctive kind of representational perspective; specifically, a rational individual is a *locus of incompatibilities* for propositional attitudes, as well as, more generally, a *subject of normative statuses*. Rational individuals do not simply have states, they can be *entitled* to these states, or *obligated* to relinquish them. It is this that constitutes an individual's distinctively *rational* perspective, and makes her a locus of distinctively *conceptual* representational states.

Furthermore, at least part of what it is for a rational individual's perspective to constitute a point of view on the world is that the attribution of states to such individuals is governed by norms of consistency and coherence, mediated by relations of inferential consequence. To be a rational individual is thus, amongst other things, to be a creature that *should not* have incompatible beliefs, or deny the deductive consequences of her beliefs.

If Benedict believes that Pangur is a cat, he *should not* believe that Pangur is a dog. Such relationships between propositional states like belief are partly what constitutes the distinctively rational unity of his representational perspective, and they are partly what make it *his* perspective: If Benedict believes that Pangur is a cat and *Anthony* believes that Pangur is a dog this generates no obligation for Benedict to revise his beliefs, because Anthony has a different representational perspective—*is a different individual*—than Benedict.

Rational creatures are also capable of basing their beliefs on other states that are part of their rational perspective, and in fact *only* on states that are part of their perspective. Benedict may base his belief that Pangur is an animal on his belief that Pangur is a cat,

³²I say '*principally* propositional attitudes' because not all conceptual content is propositional. Burge rightly notes that concepts are essentially possible elements in propositional contents (Burge, 2010, p. 540), but he wrongly concludes on this basis that concepts can occur *only* in propositional structures. McDowell's intuitional content is conceptual but non-propositional (see McDowell, 1998b, 1998c, 2008a).

but he could not base his belief that Pangur is an animal on *Anthony's* belief that Pangur is a cat. Individual rational perspectives thus constitutively constrain the rational relations, such as the relation of one belief being that individual's reason for another, that may hold between the propositional attitudes that are constituents of these perspectives.

One implication of the constraint that rational perspectives place on the rational relations that hold between their constituents is that rational perspectives consist not simply of states that bear rational relations to one another, but also of specific instances of such relations between such states. The particular relations that constitute an individual's perspective are specific instances, in that individual's particular psychological perspective, of some of the normative relations described above that govern, and help constitute the distinctive unity of, rational perspectives in general.

To illustrate, suppose that both Anthony and Benedict believe that Pangur is a cat, that Pangur is an animal, and that Pangur has four legs. If Anthony based his beliefs that Pangur is an animal and that Pangur has four legs on his belief that Pangur is a cat, he is obligated to rescind both of these beliefs if he abandons his belief that Pangur is a cat.³³ But if Benedict instead based his beliefs that Pangur is an animal and that Pangur is a cat on his belief that Pangur has four legs, he would *not* be obligated to rescind his belief that Pangur is an animal if he abandoned his belief that Pangur is a cat, for he would still have the belief that serves as his reason for this belief.³⁴

Explanations of a rational individual's perspective thus make essential appeal not only to her propositional attitudes, but also to the normative relations that hold between these states, such as relations of inferential consequence; of what counts as a reason for what. As

³³If he has some other belief that can serve as the basis for these beliefs he need not, but I set this complication aside. For more on the logic of belief revision, see Alchourrón et al., 1985.

³⁴Benedict is perhaps slightly odd, since it is not very epistemically responsible for him to base his belief that Pangur is a cat on his belief that Pangur has four legs *alone*, but this doesn't affect the point being made, since the relevant norms needn't be, and I think *aren't*, exhausted by relations of deductive consequence.

the previous example illustrates, we cannot explain some of the normative features of a rational individual's states, such as Benedict's entitlement, and Anthony's lack of entitlement, to the belief that Pangur is an animal (after abandoning the belief that Pangur is a cat), except by reference to such relations.

Furthermore, that these states have these normative features will often play an essential role in *psychological explanations*: Anthony, upon learning that Pangur is not a cat, might thereby give up his belief that Pangur has four legs precisely *because* he recognises that he is no longer entitled to it, just as he might initially have come to believe that Pangur is an animal *because* he recognised that it was a consequence of his belief that Pangur is a cat.

2.4.2 Rational Perspectives and Rational Perception

Burge has demonstrated that a creature does not *need* to have a rational perspective of this kind in order to be able to objectively perceive, or for these states to be attributable to it as an individual. Whatever the representational perspectives of non-rational creatures are, they are not rational perspectives of the kind just described, for the simple reason that such creatures do not have propositional attitudes and are not subject to epistemic or semantic norms. And yet, for reasons we saw above, they are nonetheless capable of objectively perceiving, and such perceptual states feature in the explanation of their (non-rational) perspectives.

But suppose that we *are* dealing with a creature with a perspective of this kind; a creature with propositional attitudes that stand in rational relations to one another in the way described above. What *in this case* should we say about perceptual states attributable to a creature with this kind of perspective? The answer I want to give to this question in this section is that the distinctively rational nature of a rational individual's perspective

means that perceptual states attributed to her constitutively involve rational relations to her propositional attitudes, and must at least be capable of serving as the basis for beliefs.

Intuitively we can see why this should be so because, *given* that a rational individual's perspective is constituted in part by its propositional attitudes, and the rational relations these propositional attitudes bear to one another, attributing a state to a rational individual *just is* placing it within this distinctively rational nexus. To say that *Benedict* is perceiving is to say that the individual that, amongst other things, believes that Pangur is a cat is perceiving. And the unity of Benedict's representational perspective is provided by the rational norms that relate conceptually contentful states to those other states that constitute his perspective.

Perceptual states are constitutively attributable to individuals because, in part, they are constituents of the perspectives of individuals, and serve to explain features of their perspectives. But since rational individuals have perspectives that are constituted in part by propositional attitudes, and are governed by rational norms of coherence and consistency, to say that perceptual states are constitutively attributable to such individuals is to say that perceptual states are constituents of *rational* perspectives, and feature in the explanation of features of such perspectives.

Thus, if *Benedict* sees a white cat, to attribute this perceptual state to *him*, rather than merely to his visual system, means attributing it to someone who has the propositional attitudes that Benedict has. If Benedict has the belief that all cats are black, this means that we are attributing to him a state that is *incompatible* with one of his beliefs. Benedict may thus give up his belief that all cats are black in light of his current perceptual state, and he may judge, on the basis of his current experience, that some cats are white. His perceptual state is thus capable of having the same kind of bearing on what he ought to think, and in

some cases what he *does* think, as his beliefs do. We cannot explain this, however, except by appealing to the rational features of Benedict's perspective described above.

So far we have been drawing out, in at least one plausible way, the implications of the above characterisation of the notion of a rational perspective for Burge's claim that a constitutive condition of objective perception is that the perceptual state be attributable to an individual. However, I am attempting to show not merely that this kind of intellectualist view of perception is plausible, but also that Burge's claim that perceptual states are constitutively attributable to individuals and yet constitutively independent of her rational capacities is *not*.³⁵

We can turn these observations into something more closely resembling an *argument*, however, by considering what it would be like to attribute a perceptual state to a rational individual with a rational perspective and yet for that perceptual state to *fail* to stand in rational relations or be capable of serving as a reason for belief in the way described above. To this end, consider the following example:

Benedict is very much like his cat Pangur. He has an objectifying visual system that gives rise to perceptual states that objectively represent objects in his environment and some of their visible properties, just like Pangur. In another respect, however, Benedict is very much *unlike* his cat. For Benedict is rational and thus *also* has propositional attitudes and is capable of basing his beliefs on reasons. Right now, Benedict's visual system is representing to him the visual scene that is in front of his eyes. He is in an objective perceptual state.

Although the other states attributable to Benedict stand in rational relations to his propositional attitudes, and are capable at least in principle of serving as reasons for belief, his perceptual states, though *his*, are constitutively independent of all of his rational capacities.

³⁵I'm hesitant to use a stronger term than 'plausible' because the issues here are delicate and require careful treatment. Aiming for a charge of *inconsistency* seems to be overly optimistic. Thus, although in the characterisation I give of the example below I claim that the situation described is incoherent, I remain open to the possibility that this judgement itself can be made consistent with Burge's view somehow. As we will see in the conclusion, however, modesty even here in the more ambitious argument is not necessarily a bad thing.

Thus, even though Benedict is in a perceptual state that represents the inkwell on the table in front of him, this state bears no rational relations to his other states, and he is incapable of basing beliefs on this state.

Were you to ask him, therefore, he would deny that he is seeing anything. Even if somehow he were to form a belief that in fact is true of some aspect of the visual scene in front of him, his perceptual experience could not be a reason for this belief, and he could not have formed it on the basis of this perceptual experience. And if he formed a belief that is *false* of some aspect of the visual scene in front of him, his perceptual experience would not be incompatible with this belief, and could not be the basis for him relinquishing this belief.

I do not think this is a description of a coherent psychological situation. A state of one of Benedict's subindividual systems, such as his digestive system, might play this kind of role in Benedict's life. But a state attributable to Benedict *as an individual* could not. Benedict's perceptual state is functionally independent of the propositional attitudes that constitute his representational perspective, and of the norms that constitutively govern his perspective. This is as much as to say that it is not a part of his rational perspective at all.

Importantly, Benedict's perceptual state will not even be able to play the same kind of role in his mental life as the same kind of perceptual state could play in Pangur's mental life. For perception to do for a rational individual even the relatively unsophisticated things it does for a non-rational individual like a cat, it must be attributable to Benedict as an individual. But, as we have just seen, there is no reason for us to suppose that this state *is* attributable to him in this way, given that it is not part of his rational perspective in any obvious sense. If his perceptual state could, for example, cause him to move his body in something like the way that it causes Pangur's body to move when Pangur pounces on the mouse, it would cause his body to move for reasons unknown to him, reasons inaccessible to his rational capacities. His perceptual state might cause his body to move but we could

not call the resulting bodily movement *his action*, and could not explain it by reference to states attributable *to him*.

The crucial fact, therefore, is that once an individual acquires intellectual capacities it transforms *the individual herself*. She, as an individual, is no longer the kind of creature she was before. She is now a rational individual with a distinctively *rational* perspective. States attributable to her are thus attributable to her as such. To put the point somewhat more poetically than we have done so far, she is now a *person*, rather than a *mere* individual (cf. Kant, 1798/2006, p. 15). A perceptual state cannot feature in the explanation of the perspective of a person, and hence cannot be attributed to her as an individual, unless it is capable of standing in rational relations with her propositional attitudes and of serving as a reason for belief. And so, given that featuring in such explanations is a necessary condition for a state to be objective, we cannot explain what ‘objective perception’ signifies in the case of persons without reference to their rational capacities.

2.4.3 Some Responses

I argued, in the previous section, that rational individuals have distinctively rational perspectives, and that as a result perceptual states are only attributable to rational individuals if they are capable of standing in rational relations to the conceptually contentful states of such individuals, and capable of serving as the basis for the formation, or relinquishment, of such states. In this final section I want to argue that this kind of intellectualist view is *not* a version of individual representationalism, and does not conflict either with Burge’s view of objectivity or with vision science more generally.

Although Burge does not address the specific kind of view I outlined in the previous section, his view, for reasons that by now are probably predictable, nonetheless entails a

rejection of this kind of view. Burge claims that perceptual states are constitutively independent of rational creatures' propositional attitudes, rational capacities, and of uniquely rational norms (as opposed to the norms of veridicality bestowed upon perceptual states merely by virtue of the fact that they are representational). Here is what Burge says in a discussion of Sellarsian views that claim that perception has propositional content:³⁶

“The form of perception is not a matter that submits to stipulation or to other armchair pronouncements that derive from considering epistemic norms rather than perception itself. Whether perception has propositional form depends on what perception is... There is no strong reason to believe that attribution of propositional capacities is needed in empirical explanations of the perceptual systems of bees, frogs, birds, or even humans. The issue is not advanced by echoing Sellars's claim that empirical warrant must occur within a 'space of reasons'.” (Burge, 2010, p. 434)

We are in a position to see now why this point, made by Burge at a number of different places, misses the point. I have argued that a rational individual has a rational perspective, and normative relations such as relations of inferential consequence and incompatibility govern the perspectives of rational individuals, and are instantiated in the relations between the propositional attitudes that partly constitute such perspectives in such a way that they may feature in psychological explanations. Burge's criticism fails to connect with this kind of view in at least three respects.

Firstly, the appeal to explanations of the perceptual systems of bees, frogs, and birds, presupposes an additive theory of rationality, and so begs the question against the current view. The perceptual states of rational creatures might constitutively involve rational relations that feature in the explanation of such individuals' rational perspectives even if this is not true of non-rational animals. I will not dwell on this point since we have already dis-

³⁶Although I focus only on this quotation, the points I make below apply to innumerable similar other claims that Burge makes throughout *Origins of Objectivity* in connection with many other authors, too.

cussed similar points at length above, but I want to note one aspect of this response specific to the current dialectical context before moving on.

It would not advance the dialectic if I rejected some of Burge's arguments on the basis that they presuppose an additive theory of rationality only to presuppose a transformative theory of rationality in the current argument. We are not in danger of this, however, because the argument given in the previous section only presupposes that individuals have rational perspectives, by virtue of being bearers of rational states. The distinction between transformative and additive theories of rationality concerns how these views conceive of the relationship between the rational states and capacities of rational individuals and those states and capacities they share with non-rational creatures. But the argument I made above did not presuppose anything *about perception* at all. It only presupposed this view of rational perspectives, along with Burge's own claim that perceptual states are constitutively attributable to individuals, and from these assumptions drew out the consequences for perception.

Furthermore, even Burge acknowledges that propositional attitudes are attributable to rational individuals (Burge, 2010, p. 540), that belief is essentially a certain kind of commitment governed by epistemic as well as semantic norms (Burge, 2003, p. 524), that having a propositional attitude requires the capacity to draw inferences from one such attitude to another (Burge, 2003, p. 526fn), and that rational creatures are capable of citing reasons for their beliefs (Burge, 2003, p. 435). He even says that he agrees with Sellars (and Kant) that "*reasons* that support knowledge must be propositional" (Burge, 2010, p. 435).³⁷ He simply denies that these facts are relevant for understanding the nature of the perceptual states of

³⁷Even if Burge *did* deny some of these facts, which he shouldn't, it would not *beg the question* against him, since the point under discussion is not whether rational creatures have such states and capacities but what the relation between these and their perceptual states are.

rational creatures. I have argued in effect that if, as he claims, such states are attributable to rational individuals, rather than their subsystems, these facts *are* relevant.

Secondly, Burge's focus on the relevance of propositional capacities to empirical explanation of the *perceptual systems* of humans is a red herring. We can agree that there is "no strong reason to believe that attribution of propositional capacities is needed in empirical explanations of the perceptual systems of bees, frogs, birds, or even humans". On the current account propositional capacities are relevant not for understanding how perceptual states are *produced* by the visual system but rather for understanding what is necessary for such states to be *consumed* by the individuals to which they are attributable (cf. Millikan, 1989).

Even on Burge's account the explanations relevant for determining the nature of perceptual states are not (or at least not *solely*) explanations of a creature's *perceptual system* but rather the role that perceptual states themselves play in the psychological explanations of individual perspectives and actions. Burge's criticism may apply to some of the Sellarsian accounts that are his explicit target in the quotation above, but they do not bear on the current view, which is motivated by considerations concerning the role that perceptual states play in the explanation of a rational individual's perspective, by direct analogy with Burge's own arguments for his transformative conception of objectivity.

Thirdly, and finally, Burge puts the issue here, as he does often elsewhere, in terms of a conflict between empirical considerations deduced from vision science and epistemic considerations, which he claims are irrelevant for understanding the nature of perception. This characterisation of the disagreement is misleading not only because it repeats the familiar narrative about a clash with science that I have debunked above, but also because it misunderstands the relevance of appeals to rational norms, such as that of one state being a reason for another, that belong to the 'space of reasons'.

Although the relevant inferential relations, and the norms of consistency and coherence, that characterise a rational individual's perspective are epistemic in the sense that they determine the rational norms to which that individual is subject when she reasons or forms her beliefs, the explanations given above are not for that reason epistemic. Rather, they are principally of *semantic* and *psychological* significance: They feature in explanations of *what* an individual believes—what kinds of experiences *entail* this belief and what kinds of experiences are *incompatible* with it—and *why* they believe it—e.g. because their experience gives them a *reason* for that belief.

In providing the characterisation of a rational individual's perspective I did not even attempt to explain the conditions under which she would be *justified*, or have *good* reason for her beliefs, as I would have if I were attempting to motivate this view by means of epistemic considerations, and as Burge suggests some Sellarsians do. Rather, the appeal to notions like relations of inferential consequence and inconsistency was invoked principally as a way of explaining what, psychologically, an individual's perspective is and what kinds of relational semantic properties, such as inconsistency and entailment, the states that constitute her perspective have. The need for perceptual states to have a place in the 'space of reasons' is thus not, in my argument at least, motivated by a conception of warrant but rather by a conception of the notion of a rational perspective.

There is obviously more that could be said about this point, given the sheer breadth of Burge's book, as well as the treatments of similar issues he has given elsewhere. But the three responses just given cover the central objections that Burge has to intellectualist views that give a role to propositional attitudes in their explanation of perception, which are the objections most relevant to the account given in the previous section.

We are now in a position to explain *why* it is that our account of objective perception

for rational creatures must make reference to those creatures' rational capacities even if this is not true of non-rational animals, and the sense in which perception is *transformed* by the acquisition of rational capacities. Having an objectifying visual system is, as we saw, a necessary but not sufficient condition for objective perception. Objective perception in rational creatures constitutively depends on their rational capacities not because this is required for objective perception *tout court* but simply because objective representational states have to be attributable to them *as individuals*. Objective perception in non-rational creatures does not constitutively depend on rational capacities because such creatures do not have rational perspectives.

The capacities that distinguish rational creatures from non-rational ones thus allow the former to do much more than non-rational creatures can, but this increase in ability comes at a cost. For it means that for a state to be attributable to a rational creature at all, she must be able to do *more* with it than a non-rational creature can. It must be integrated into her rational perspective, governed by the same rational norms that govern her propositional attitudes, and it must be capable of serving as a reason for such attitudes. For a non-rational creature, by contrast, all that is required for a state to be attributable to it is that it can be integrated into her *non-rational* perspective. This view thus does not deny that non-rational creatures or infants are capable of objectively representing, it simply insists that *once* a creature has rational capacities its perceptual states must be properly integrated with such capacities.

To end this section, consider an analogy. Suppose that a species of frog is capable of leaping two feet in distance. And suppose that over time—in response to increasingly sophisticated predators, perhaps—it evolves a number of new muscles that allow it to leap *four* feet. Clearly when the strong frog leaps four feet it uses muscles unavailable to its

more primitive ancestor. And this new, more sophisticated creature can, obviously, also still leap two feet in distance. But when it *does* leap two feet it does not, indeed *could* not, do so simply by activating only those muscles it shares with its less sophisticated ancestor. Rather, to do even what its less sophisticated ancestor could do—leaping two feet—requires the creature to do more than its less sophisticated ancestor could do—namely, to activate precisely those muscles that differentiate it from this ancestor.

Acquiring rational capacities, like acquiring extra muscles, gives us the capacity to do things that our less cognitively sophisticated ancestors, or even ourselves as infants, could not do. But, also like acquiring extra muscles, these new aspects of our mindedness become interwoven with the old and cannot be separated from them. To do what our ancestors could do with less sophisticated abilities, or what we could do as infants before our rational capacities fully developed, requires us to draw on the capacities we share with them—we could not perceive unless we had objectifying visual capacities, after all—but *in combination* with those that are unique to us, and allow us to also do much more than them. When we perceive, we are like the strong frog who only leaps two feet: We are doing something that simpler creatures than us can do, but precisely because we are more developed, we are doing it in a way that draws on those capacities that are unique to us.

2.5 Conclusion

If the argument I have given in this chapter is correct, intellectualism about perception does not conflict with visual psychology, or with Burge's account of the constitutive conditions for objective representation in non-human animals and infants. Burge frequently claims that armchair considerations regarding epistemology cannot take precedence over the deliv-

erances of common sense and the psychology of vision; that epistemology “cannot dictate to psychology” (Burge, 2010, p. 435).

Were it true that intellectualism is simply incompatible with our best scientific understanding of vision it would indeed be a devastating indictment of this approach. But if the discussion above is correct, the central disagreement between Burge and intellectualists—at least once we separate the latter from the pessimistic view of non-rational perception that I have argued is optional—rests principally on epistemic, semantic, and metaphysical questions whose answers science alone does not dictate. Only by assuming answers to some of these questions is there a potential for conflict with the results of vision science.

In a sense, however, even the more ambitious aspects of the argument I have been making are still relatively modest. After all, there are still a number of significant points of disagreement between Burge and intellectualists that I have not resolved by the discussion above, even where I have briefly touched on them. What matters is that they are questions internal to philosophy rather than being clashes between philosophy and science.

Despite the relative modesty of this argument, however, its implications are profound, for it relocates the disagreement between intellectualists and their empirical detractors to ground much less obviously favourable to the latter. The claim that intellectualists have a mistaken conception of warrant, or that they overintellectualise semantics, is much less forceful than the claim that they are ignorant of science, or that their view depends on a mischaracterisation of the perceptual capacities of non-rational creatures.

Finally, now that we have seen the view I have defended in detail, I want to return to the issue raised right at the beginning of the introduction to this chapter of the decision to choose Burge as the principal representative of empirical criticisms of intellectualism. The main positive argument of the chapter given in §2.4.1 clearly depends on the idea

that perceptual states are constitutively attributable to individuals. If they were not, the argument that attributing a state to an individual entails that it must stand in rational relations to propositional attitudes would be irrelevant for demonstrating anything about the nature of perception.

Thus, the argument I have given here will not gain any traction with the arguments of empirically-minded philosophers of perception who do not accept this aspect of Burge's view, and who think that perceptual states can be subindividual. I have two points to make in response to this observation. Firstly, I think Burge's arguments against such accounts are good, possibly even decisive. This is in part why I gave them in such detail in the relevant sections. They not only set up the argument I give but, in the context of thinking about the response to take to other empirical philosophers of perception, can be considered part of this argument.

Secondly, if these arguments are *not* good then whilst this has the unfortunate implication that the current argument for intellectualism would have to be abandoned, or at least significantly altered, this would not spell the end of intellectualism more generally, for it would also have the implication that Burge's arguments against intellectualist views would not hold. The argument I have given here therefore incorporates parts of Burge's own argument at least in part because if Burge's arguments are right, intellectualists are precluded from making many of the arguments they typically give for their view, and so are forced to search for new ones. But if Burge's arguments are wrong, they are back on the table.

Burge's arguments against individual representationalism, and his positive conception of objectivity have been taken, I think rightly, to change the state of play in the philosophy of perception. For example, Matt Nudds says that the book is "the most systematic account of perception and perceptual content... to have been published for decades" (Nudds, 2012,

p. 157), and Christopher Peacocke calls it “a landmark of contemporary philosophy” that “will be studied for years to come” (Peacocke, 2014, p. 477). If the argument I have given here is correct then even if Burge’s argument spells the death of individual representationalism, it need not spell the death of intellectualist approaches to perception more generally.

3.0 An Inferentialist Account of Perception

“What entitlements I have is not settled by what entitlements I exploit... Perceptions, as I see them, are experientially acquired entitlements to belief... This entitling circumstance is, by all means, a propositional attitude... But it is not itself a belief... [T]hat the entitling circumstance itself needs to be distinguished from that belief is brought out by the fact that one can have the entitlement without realizing that one does, and so without having the belief at all.”

—John McDowell, ‘Reply to Kathrin Glüer’, p.215

3.1 Introduction

Inferential semantics originated with inferential approaches to logical constants. This kind of approach arguably began with Frege’s *Begriffsschrift* (1879), but certainly it took root in earnest with Gentzen’s definition of the meaning of logical constants in terms of rules for their introduction and elimination (Gentzen, 1934a, 1934b). From here, via Sellars (1953, 1974), Dummett (1975, 1979), and Harman (1975, 1982, 1987), similar approaches were taken to explain the meaning of *nonlogical* terms. Since then, largely under the influence of Robert Brandom, it has evolved into a wide-ranging research programme centred around the philosophy of logic and the philosophy of language, but with implications for the philosophy of mind (Fodor & Lepore, 2001), epistemology (Brandom, 1995, 2002a; McDowell, 1995), metaethics (Warren, 2015), the philosophy of science (Khalifa et al., 2018), the history of philosophy (Brandom, 2002b, 2019), and even pedagogy (Radford, 2017).

Despite this wide-reaching influence, inferentialism, at its heart, is characterised by a constellation of commitments in the philosophy of language. In this chapter I explore the

connections between the core commitments of inferentialism as an approach to language and spell out what I think the implications of these might be for the philosophy of perception. My hope is to isolate a small part of the inferentialist view—roughly the part laid out in the first three chapters of *Making it Explicit* (Brandom, 1994a)—and to show both how the inferentialist view of language benefits from taking perception more seriously and how the resulting view is of interest as a view of perception in its own right.

The core part of the inferentialist view of language on which I will focus consists of an *inferential semantics* grounded in a *normative pragmatics*. There is often disagreement about whether to classify use theories of meaning such as inferentialism as semantic theories (e.g. Speaks, 2019) or rather as metasemantic theories (e.g. Burgess & Sherman, 2014a, pp. 10–11). Below I explain inferential semantics and its relationship to normative pragmatics in more detail, but as a background to this explanation it will be helpful if we provide a general characterisation of the distinction between semantics and metasemantics here first.

The distinction between semantics and metasemantics originates with Kaplan, who introduces it as follows:

“The fact that a word or phrase has a certain meaning clearly belongs to semantics. On the other hand, a claim about the *basis* for ascribing a certain meaning to a word or phrase does not belong to semantics... Perhaps, because it relates to how the language is *used*, it should be categorized as part of the *pragmatics*... or perhaps, because it is a fact *about* semantics, as part of the *Metasemantics*... (or perhaps, for those who prefer working from below to working from above, as part of the *Foundations of semantics*)” (Kaplan, 1989, pp. 573–574)

A *semantic theory* is thus a theory of the meanings of the expressions—words, phrases, sentences, etc.—of a language. A *metasemantic theory*, by contrast, is a theory of how, or by virtue of what, expressions come to have the meanings they do.¹

¹The distinction between semantics and metasemantics traces back to debates that arose out of Kripke’s

In practice, semantic theories typically consist of explanations of the meanings of *classes* of expressions, although particularly troublesome individual words or constructions sometimes warrant special attention of their own. Semantic theories (especially formal ones) also often spell out principles that explain how the meanings of phrases and sentences depend on the meanings of words and their mode of composition. Inferentialism, as we shall see, characterises the meaning of sentences in terms of their role in inference, and the meaning of subsentential expressions in terms of the contribution they make to the inferential roles of sentences in which they feature.

Metasemantics is both less well developed and more diverse as a discipline than semantics, but in practice most metasemantic theories consist of accounts of the *metaphysics* of meaning (cf. Burgess & Sherman, 2014b). However, as Kaplan hints in the quotation above, they may, either as part of or instead of giving a metaphysical account, include an explanation of the *pragmatics* of the language; how expressions are *used* by speakers of that language. A pragmatic theory is not, in itself, a metasemantic theory, but it can be put to use as one, if it can be shown how the meanings of expressions are grounded in their use. The inferentialist metasemantics is a theory of this kind: It treats the inferentialist account of pragmatics as the basis for an explanation of how sentences of a language come to have the meanings they do.²

causal theory of reference. Aside from Kaplan, who identifies it as the central point at issue in the debate between Kripke and Donnellan (Kaplan, 1989, p. 573fn), Stalnaker also draws the distinction (using the terminology of a ‘foundational theory’) in a paper on Kripke’s theory of names (Stalnaker, 1997, pp. 166–7). However, Burgess and Sherman identify Dummett’s notion of a theory of meaning (Dummett, 1976) as (an element of) metasemantics (Burgess & Sherman, 2014a, pp. 2, 10–11), so the distinction also has a good inferentialist pedigree (though it is not clear that Dummett’s tripartite distinction between a theory of reference, a theory of sense, and a theory of force lines up neatly with the distinctions drawn by contemporary semanticists and metasemanticists).

²This task concerns the first two of three stages in Brandom’s account (cf. Wanderer, 2008, p. 97). He then proceeds to provide an explanation of the meaning of subsentential expressions, and of representational notions like truth, in part 2 of *Making it Explicit*. I talk about one element of this part of the project briefly in §3.5.2.

In the philosophy of mind there is an analogous distinction between first- and second-order questions about the content of *mental* states. A central question in the philosophy of mind is the question of how, or by virtue of what, mental states such as beliefs or perceptual experiences come to have content. As this way of putting the question makes clear, there is a close analogy between such *foundational theories of mental content* and metasemantic theories.³ Indeed, there are sometimes close connections between foundational theories of mental content and foundational theories of meaning. The most common foundational theories of content are explicitly naturalistic and usually reductive, grounding the content of mental states in naturalistically respectable notions like proper function or information (cf. Burge, 2010; Dretske, 1981; Fodor, 1987; Millikan, 1989; Papineau, 1984), and some of these are used as the basis for an account of meaning or linguistic content (Millikan, 1984).

In addition to this second-order question regarding how mental states come to have content, however, there are also a cluster of first-order questions concerning the nature of the content of mental states themselves. Such questions are particularly central to the philosophy of perception, given that the nature of the content of perception is less clear than in the case of propositional attitudes like belief. Answers to such questions often vary significantly along a number of important dimensions, some of which we will touch on in the next chapter. For now I note these questions primarily to set them aside.

In this chapter I focus instead primarily on the second-order question of how, or by virtue of what, perceptual experiences come to have content. In the previous chapter, I argued that for a perceptual state to be attributable to a rational individual, it must bear rational relations to that individual's propositional attitudes, and must be capable of serving as a

³The use of the term 'foundational' is meant to invoke Kaplan's use of the phrase 'foundations of semantics'. Given that the term 'metasemantics' applies at best only awkwardly to mental content, I will use the phrase 'foundational theory of (mental/perceptual) content' when I want to highlight the analogy between the second-order projects in the philosophy of language and the philosophy of mind.

reason for belief. In this chapter, I provide an inferentialist account of the foundations of perceptual content that grounds the latter in this normative, reason-giving role of experience, and explain how this relates to the constellation of views that form the core of the inferentialist approach to language.

As we will see, Brandom has his own explanation of the role of perception in inferentialist semantics.⁴ However, I argue below that Brandom's account is unsatisfactory, both for reasons relating to the philosophy of language and for reasons more directly relating to the philosophy of perception. By detaching Brandom's metasemantic framework from his more idiosyncratic view of perception, I aim to provide a more plausible account of perceptual content from within this inferentialist framework.

3.2 Inferentialism

3.2.1 Inferentialist Semantics

The *central semantic thesis of inferentialism* is the claim that the meaning of a sentence is its inferential role; the role that the sentence plays as a premise and conclusion in good inferences (Peregrin, 2006, p. 1; Peregrin, 2014, p. 3; Brandom, 2000, p. 11). Inferentialism is thus a *normative use theory* of meaning. The view counts as a use theory of meaning because it conceives of the meaning of a sentence in terms of features of its use; specifically, in terms of the role it plays as a premise or conclusion of inferences. And the view counts as an explicitly *normative* version of the use theory because the inferences

⁴I say that Brandom has an inferentialist approach to *perception* but as we will see this is somewhat misleading, as his approach vitiates the need for an account of perception as such, by replacing it with an account of perceptual *judgement* in which experience plays at best a merely causal role. I will continue to speak this way, however, for the simple reason that there is no other obvious alternative that isn't too awkward or unwieldy for regular use.

that constitute the meaning of a sentence are the *good* ones—those that a speaker *ought* to endorse, which need not be coextensive with the inferences that she *does* endorse.⁵

It will help, for the purposes of highlighting the unique features of the view, to contrast inferentialism with two other (classes of) views: *Representational* semantics and *descriptive* use-theories of meaning.

Use theories of meaning are usually contrasted with *representationalist* theories of meaning, which characterise meaning in terms of representational relations such as reference and truth. The majority of representational theories of meaning are *correspondence* theories, which claim that “meaning is a relation between the symbols of language and certain entities which are independent of language” (Gamut, 1991, p. 1). Paradigmatic examples of such theories are *referential* theories of meaning such as Kripke’s Millian theory of proper names: On Kripke’s theory the meaning of a name just is the object that name (rigidly) designates (Kripke, 1980).⁶ Although straightforward correspondence theories of this kind are most obviously well-suited to the semantics of names, it is possible to extend this approach—say, by associating properties with predicates—to build a general semantic theory of this kind.

However, to be a representationalist one need not think of meanings as *things* associated with expressions, or of the job of a semantic theory as that of associating such entities with expressions of the language. Davidson, for example, suggests that a theory of truth can do the work of a semantic theory for natural languages (Davidson, 1965, 1967, 1970). On a *truth-conditional* account of this kind, a formal theory of truth can serve as a theory of meaning for a language provided it entails, as theorems, statements of the truth-conditions of sentences

⁵The same point applies at the level of entire linguistic communities, too. Cf. Brandom, 1994a, ch.1.

⁶Kripke combines this view with a *causal theory* of names, which is better thought of as a *metasemantic* claim about how it is that names come to have the meanings (referents) they do (cf. Stalnaker, 1997, p. 167).

of that language. Although simple referential theories and truth-conditional theories are by no means exhaustive of representational approaches to semantics, the majority of the most commonly used semantic theories fall into one of these two camps, or combine features of both, so this sketch will suffice for drawing out the contrast with inferential semantics.⁷

Inferentialism, by contrast with representational approaches, claims that the meaning of a sentence is the role that it plays as a premise and conclusion in good inferences. Thus, on a Davidsonian truth-conditional account the meaning of the sentence “Pangur is a cat” is given by the Tarski biconditional “‘Pangur is a cat’ is true iff Pangur is a cat”. On an inferentialist account, by contrast, the meaning of “Pangur is a cat” is its inferential role, which consists of the good inferences in which it features; both those—such as the inference from “Pangur is a white cat” to “Pangur is a cat”—in which the sentence features as a conclusion, and those—such as the inference from “Pangur is a cat” to “Pangur is an animal”—in which the sentence features as a premise.

It is crucial for the plausibility of inferentialist semantics that the inferential role of a sentence includes not just *logically* good inferences, but also *materially* good inferences. Thus, Sellars emphasises that the inference from “Pangur is a cat” to “Pangur is an mammal” is good in its own right, even without the collateral premise “All cats are mammals”. And according to Sellars this is of some semantic significance, for it is principally material, rather

⁷Most contemporary semanticists are formal semanticists, and most formal semanticists work broadly in the tradition of Montague semantics (see, e.g. Gamut, 1991). Such theories ordinarily supply a model that associates (typically abstract) entities with expressions of the language, and so are to this extent correspondence theories. However, such theories generate specifications of the truth-conditions of sentences as theorems, and so to this extent appear to serve the general aims of a truth-conditional theory. The relationship between such theories and Davidson’s project are complicated, however, because formal theories typically supply truth-conditions in terms of a relativised notion of truth such as truth in a model or truth at a world, rather than the absolute notion of truth that is required to satisfy Convention T. Thus, although some formal theories entail instances of the T-schema properly so called (e.g. Heim & Kratzer, 1998, chs.1–2), Davidson himself insisted on distinguishing his view from theories that only provide specifications of truth-conditions relative to a model, world, or interpretation. For Davidson’s discussion of this issue, see Davidson, 1973a, especially pp. 68–70, and for a suggestion regarding *exactly* what Davidson thinks the problem with such theories is, see Joseph, 2004, pp. 204–5.

than logical, inferences that articulate the conceptual content of a sentence and give it its meaning. Thus, although “Pangur is a cat” implies both “Pangur is an mammal” and “Something is a cat”, the first of these inferences tells us more about the meaning of the sentence than the second, for it tells us about the conceptual connections between concepts that give the term “cat” its *distinctive* content. As Sellars puts it:

“[Material] transformation rules determine the descriptive meaning of the expressions of a language within the framework established by its logical transformation rules... In traditional language, the “content” of concepts as well as their logical “form” is determined by rules of the Understanding.” (Sellars, 1953, p. 25)

This point highlights an important explanatory difference between inferentialism and *referential* representational approaches. The fact that the inference from, e.g., “Pangur is a cat” to “Pangur is a mammal” is a good one is taken as semantically basic for the inferentialist.⁸ From the perspective of representational semantics, this puts the inferential cart before the referential horse: On such accounts the good inferences a sentence features in are *consequences* of its representational properties. That is, in the specific case in question, a representational semanticist would likely insist that it is *because* the extension of the term “cat” is a subset of the extension of the term “mammal” that “Pangur is a cat” entails “Pangur is a mammal”. For the inferentialist, by contrast, the good inferences a sentence features in are not *consequences* of its meaning; they are *constitutive* of it:

“meanings do not determine the appropriate inferences, but what inferences are socially appropriate determines the meanings of the sentences involved in those inferences” (Brandom, 1983, p. 644).

The claim that the meaning of a sentence is the role it plays in inferences, however, admits of at least two interpretations. On the first, *descriptive*, interpretation, the rel-

⁸Though as we will see in §3.4.1 below, it is grounded in certain *pragmatic* facts.

evant processes of reasoning are those patterns of reasoning that speakers of a language *in fact* engage in (or at least are disposed to engage in). This, for example, is how Horwich characterises the relevant features of use (cf. Horwich, 2005, pp. 28, 32):

“For each term there is a small set of simple properties which (in conjunction with other factors and with the basic properties of other words) explain total linguistic behaviour with respect to that word... The distinctive form of [the semantic] feature [of use of words] is that it designates the circumstances in which certain specified sentences containing the word are accepted; and the primary explanatory role of a word’s acceptance property is to account for the acceptance of other sentences containing the word.” (Horwich, 1998, pp. 44–45)

Horwich’s use theory of meaning is thus *descriptive* rather than normative. It is a use theory because he identifies the meaning of a term (and, by extension, sentences) with a feature of its use rather than with, say, its contribution to the truth conditions of sentences in which it features. But the relevant features are, for Horwich, the *actual* patterns of acceptance and rejection speakers engage in, rather than those they *ought* to engage in.⁹

Brandom, by contrast, says in the quotation above that the meaning of a sentence is determined by the inferences that are socially *appropriate*. We can thus contrast descriptive use theories such as Horwich’s with inferentialism, which, as a *normative* use theory of meaning, conceives of the relevant features of use as the inferences that a speaker *ought* to accept. The inference from “Pangur is a cat” to “Pangur has retractable claws” is a good one, and so partly constitutes the meaning of that sentence, whether or not the person using that sentence is actually disposed to accept the latter sentence on the basis of the former one or not.

It is difficult to characterise the difference between descriptive and normative use theories of meaning in purely semantic terms, however. Rather, the difference is more appropriately

⁹For Horwich’s discussion of normative approaches to truth and meaning, see Horwich, 1998, ch.6 and Horwich, 2005, ch.5. In addition to Horwich, see Block, 1986; Harman, 1974, 1982, 1987 for more psychologically oriented approaches.

characterised as a disagreement about *metasemantics*; in terms of a difference in the explanation of how it is that expressions come to have the meanings that they do.¹⁰ (This is plausibly why Brandom says that the meaning of a sentence is *determined* by what inferences are socially appropriate.) On this way of characterising the difference, the fundamental disagreement concerns whether to ground a use theory of meaning in practices of *actual* use or in rules for *proper* use. I thus postpone the full explanation of the normative aspect of inferentialism until §3.4.1 below, where I give the inferentialist metasemantics in more detail.

This covers the core of inferentialist semantics, but before we move on to the next section, there is one more issue we need to address, given that the focus of this paper is on issues in the philosophy of mind as well as issues in the philosophy of language. Specifically, we need to address the implications of inferentialism about meaning for the content of mental states such as belief. Brandom advocates what he calls a “*relational linguistic* approach to the conceptual” (Brandom, 2000, p. 6, my emphasis), which he characterises as follows:

“Concept use is treated as an essentially linguistic affair. Claiming and believing are two sides of one coin—not in the sense that every belief must be asserted nor that every assertion must express a belief, but in the sense that neither the activity of believing nor that of asserting can be made sense of independently of the other, and that their conceptual contents are essentially, and not just accidentally, capable of being the contents indifferently of both claims and beliefs.” (Brandom, 2000, p. 6)

One way to spell this idea out, without delving too deep into the relevant issues, is as follows. For creatures capable of using language, believing that P involves *at least* a

¹⁰I say it is *more appropriate* because I do not want to rule out that there may be some important semantic differences between the two kinds of use theory. In particular, jumping ahead a little, the distinction between commitment-preserving and entitlement-preserving inferences is, at least on its face, a semantic one; a difference in the inferential role of a sentence. It *may* be possible to spell out the difference between these two kinds of inferences in terms that don’t make reference to deontic notions like commitment and entitlement—say, purely as a difference in structural features like monotonicity—in which case we can treat such normative notions as purely part of the metasemantics. But if not we would have to give them a semantic role, which would make the difference between normative and descriptive use theories at least partly semantic, too. I take no stance on this question here.

disposition to assert that P under certain conditions (even if it involves much more than this, too). So the activity of believing cannot be made sense of independently of that of asserting. But, conversely, merely uttering the sentence “P” is not sufficient for *asserting* that P. Asserting that P is expressing a claim one’s commitment to which is not exhausted by one’s utterance of the relevant sentence on that particular occasion. Intuitively, asserting is *expressing a thought*. But more minimally, one’s assertion at least ought to be a manifestation of a general *disposition* to utter this sentence under certain circumstances (even if it involves much more than this). So the activity of asserting cannot be made sense of independently of that of believing. This is obviously not a fully developed story, but it indicates why one might think there is a non-reductive but constitutive connection between believing and asserting.¹¹

The most prominent alternatives to linguistic approaches to the conceptual are *mentalistic* approaches, which take mental content to be intelligible independently of linguistic meaning, and hence give the former explanatory priority over the latter. On Ned Block’s view, for example, the content of mental states is determined by the conceptual role of expressions in a language of thought; specifically, their *causal* role in processes of reasoning and deliberation (Block, 1986, p. 628). The meaning of expressions in a natural language is also a matter of their conceptual role, but the conceptual role of natural language expressions is “inherited” from those of this internal language of thought (Block, 1986, p. 664).¹²

¹¹For a more extended discussion of these issues, see Sellars, 1969, and for an argument for the claim that thought depends on speech see Davidson, 1975. There are also a number of much more substantive conceptions of belief and assertion that constitutively link the two. For example, most accounts of the constitutive norms of assertion entail that there is a constitutive connection between belief and assertion (e.g. Williamson, 2000, ch.11). The explanation I have given here merely provides a sketch of some plausible *necessary* conditions for belief and assertion. It is compatible with the sufficient conditions involving much more substantive constraints. For reasons already discussed in chapter 2, these are only necessary conditions for belief in rational animals. If non-rational animals are in some sense capable of having beliefs—say, states that consist solely in dispositions to act—then these conditions will not apply.

¹²Block’s view is slightly more complicated than I let on here, since he claims that meaning is constituted by *both* the internal and the external factors just mentioned (Block, 1986, p. 635). But this fact does not

The fact that inferentialism endorses a linguistic approach to the conceptual will be important in what follows for two reasons. A mentalistic approach to the conceptual would presumably require providing an account of the content of perception *prior to* and independently of the content of language. So, philosophically speaking, adopting a linguistic approach to the conceptual is important because it makes possible the kind of account of perception that I give below, which takes the content of language and perception to be interdependent. The content of sentences, I will argue, depends on their inferential connections to experiences, and vice versa.

There is, furthermore, also a practical benefit; namely, it allows us to speak of *concepts* and so justifies the use of the helpful terminology of *applying a concept*. Given that conceptual contents can be “indifferently” both contents of assertions and beliefs, one can apply the concept cat by asserting, e.g., that Pangur is a cat *or* by forming the belief that Pangur is a cat, for example by *judging* that Pangur is a cat.¹³ Most of the time it does not matter for our purposes whether a speaker applies a concept by asserting or judging, so this terminology will be useful in what follows. Similar considerations apply with respect to moving between talk of *asserting sentences* and *making claims*.

In addition, a further benefit derives from the fact that one can apply the concept cat not just in a variety of different acts but also by means of acts with a variety of different *contents*. Asserting that Pangur is a cat is applying the concept cat, but so is asserting *that cats are furry*, or *that the white cat is purring*. The only requirement for applying a concept is that it feature in an act whose content is partly articulated by that concept. Speaking of ‘applying a concept’ thus allows us to treat a number of otherwise quite different *contents*

undermine the characterisation of Block’s view as mentalistic (or at least *more* mentalistic than Brandom’s!), given his understanding of the order of priority between these two factors. The theory of mental content thus has a priority over that of the semantics of language.

¹³I will follow Brandom in typographically marking concepts by underlining words. I am also assuming that judging is one way of forming a belief.

concurrently, too, which will prove to be useful, especially in the following chapter. In a slogan: “Grasp of a concept is mastery of a use of a word” (Brandom, 2015, p. 43).

3.2.2 What *Are* Premises and Conclusions, Anyway?

In the outline of inferentialism given above the idea of something functioning as a premise or a conclusion of an inference plays a central role. For the purposes of discussing the semantics of natural language, we can treat these premises and conclusions as sentences, as we have above, or more generally, given the relational linguistic approach to the conceptual, as claims expressed by assertions of those sentences. The view that the meaning of a sentence is constituted by its role as a premise and conclusion in intra-linguistic inferences *alone* Brandom calls *hyperinferentialism*. Brandom claims, I think plausibly, that whilst hyperinferentialism might be plausible for certain terms, such as logical connectives, it is *implausible* as a general account of (natural) linguistic meaning:

“The meaning of ‘red’ is not just a matter of what other concepts its applicability is properly inferable from (e.g. scarlet), and the applicability of what other concepts are properly inferable from its applicability (e.g. colored)... Its noninferential applicability to red things is also an essential propriety of the use of the concept red: one that must accordingly be underwritten by any adequate account of the meaning or content expressed by the use of the word ‘red’. If taking the meanings expressed by logical vocabulary as a model for the semantics of other sorts of expressions could take the form only of commitment to hyperinferentialism, then there would be little reason to take this explanatory strategy seriously.” (Brandom, 2010, pp. 163–164)

It is not possible to take hyperinferentialism seriously as an approach to *natural* language because natural languages have *empirical* content, and we cannot make sense of the idea that sentences are empirically contentful if they only bear semantic relations to other sentences. Representational accounts can easily explain what makes empirically contentful expressions *empirically* contentful; namely, the semantic connections between words and the

world that are built into the foundations of the semantics itself. This is easiest to see on referential approaches, where in the simplest cases meanings just are the things denoted by expressions.¹⁴ Expressions have empirical content because their meanings, in the simplest case, just *are* mind-independent entities. Clearly *this* kind of approach is not available to inferentialists. Given that, on an inferentialist account, the semantically-constitutive relations are *inferential* relations, this means that to avoid hyperinferentialism, and make sense of empirical content, we have to give an *inferential* role to observational circumstances of application.

If the meaning of a term like ‘red’ is constituted by its inferential role, as inferentialism claims, then in order to avoid hyperinferentialism we have to expand our conception of the inferential role of such terms to include not just relations between the sentences in which it features but to also encompass the *nonlinguistic* circumstances and consequences of application of the term. Thus, in Brandom’s complete story, the inferential role of a sentence will include not just the *sentences* from which it follows and that follow from it, but also the *observational circumstances* in which it is appropriate to assert the sentence, as well as the *practical responses* that are made appropriate by an assertion of the sentence.

Thus, “This is a cat” follows from the sentence “This is a white cat”, but it also ‘follows’ in a broader sense from the visible presence of a cat. And the sentence “This is an animal” is a consequence of “This is a cat”, but if one has the desire to pet a cat, so is the action of petting it. Thus, on this view, which Brandom calls *strong inferentialism*:

¹⁴Sometimes the relationship might be more indirect, as when the meaning of a predicate is modelled as a *set* of objects, or as a *function* from certain parameters to sets of objects, but broadly speaking the connection is no mystery. Davidson’s view is, as always, more subtle. In earlier papers, Davidson claimed that relations between language and the world have a role in a truth-conditional theory of meaning by virtue of the role played by satisfaction in a Tarskian theory of truth (Davidson, 1969, p. 48). Ultimately, however, he recanted this claim in favour of an explanation of (objective) contentfulness grounded in interpretation (Davidson, 1973b, 1988).

“the visible presence of red things warrants the applicability of the concept red—not as the conclusion of an *inference*, but *observationally*. And the point is that the *connection* between those circumstances of application and whatever consequences of application the concept may have can be understood to be inferential in a broad sense, even when the items connected are not themselves sentential.” (Brandom, 2010, p. 164)

Brandom emphasises that it is the *connection* that is inferential ‘in a broad sense’ because applying a concept in response to the visible presence of, e.g., a red thing—what Sellars calls a “language entry transition” (Sellars, 1954, p. 44)—should not be thought of as *making an inference*. To form a belief that there is a cat in front of one as a result of observation is not to *infer* that there is a cat in front of one (*pace* Harman, 1973, ch.11), as one might if one formed the belief that there was a cat in front of one on the basis of one’s *belief* that one was seeing a cat.¹⁵

Rather, the transition from being in observational circumstances that warrant the applicability of a concept to the actual application of that concept is effected by means of what Brandom calls *reliable differential responsive dispositions* (RDRDs). *Perceptual* RDRDs are, as the nomenclature indicates, dispositions to respond differentially to the visible presence of certain features in one’s environment.¹⁶ In the simplest case, one might have a reliable disposition to assert “That is a cat” whenever there is a cat visibly present to one (cf. Sellars, 1968, ch.3). Ultimately, of course, this account is too simple. We do not simply blurt out “That is a cat!” whenever we see a cat in our vicinity (unless it is a particularly cute cat). Rather, more plausibly, our actual perceptual RDRDs are reliable dispositions to *judge*, or form a belief.

¹⁵The move from a belief to its practical consequences of application might count as a *practical inference*, but given that I focus on perception I set these language *exit* transitions aside.

¹⁶It does not matter for present purposes exactly what the metaphysical status of the environmental features that RDRDs respond to is. One may think of RDRDs as reliably tracking facts or states of affairs, but one might also think of them as reliably responding to events or to the presence of certain properties. Nothing hangs on this for present purposes.

Even if applying a concept observationally is not *making an inference*, nonetheless, Brandom claims, the *connection* an RDRD establishes between circumstances of application and whatever inferential consequences follow from the application of that concept deserves to be called inferential in a broad sense. Brandom explains why by considering what it means to *attribute* reliability to someone who exercises an RDRD:¹⁷

“The key point to understanding reliability as a warrantive standing in the space of reasons is that the notion of reliability itself is essentially an *inferential* notion: a matter precisely of what is a reason for what... For me to take you to be a reliable reporter of lighted candles in darkened rooms is just for me to endorse a particular pattern of reasoning; in particular it is for me to endorse the inference that could be made explicit by saying: ‘*If* in a darkened room *S* noninferentially acquires the belief that there is a lighted candle, then (probably) there is a lighted candle there.’” (Brandom, 1995, p. 905)

In other words, when one speaker treats another as a reliable reporter, what they are doing is endorsing an inference *from* the speaker’s asserting what they did in the circumstances that they were in *to* (the propriety of endorsing) the claim they asserted.¹⁸ That Benedict *said* that Pangur is a cat is, if Benedict is a reliable reporter of such things, reason for thinking that Pangur *is* a cat. But what it means for Benedict to be a *reliable* reporter is that he usually only perceptually reports that something is a cat *in circumstances in which there really is a cat visibly present to him*.¹⁹ In this way even if exercising an RDRD is not *making an inference*, RDRDs nonetheless give observational circumstances a role in good

¹⁷The move to discussing *attributions* of reliability might seem to be a *non sequitur* and, to some extent, I am sympathetic to this charge: See the discussion in §3.3.1 below.

¹⁸In the more precise terminology we will introduce later, reliability inferences are inferences from attributions of *commitment* to attributions of *entitlement* (Brandom, 1994a, p. 189).

¹⁹In *Making it Explicit*, Brandom often uses the term ‘circumstances’ to refer to the conditions under which a capacity is reliable (e.g. Brandom, 1994a, pp. 222–3). For example, my capacity to recognise colours might be reliable in some viewing circumstances, e.g. in daylight, but not others, e.g. in fluorescent light. Clearly the relevance sense of ‘circumstances’ must include *more* than circumstances of this kind if ‘circumstances’ is to mean ‘circumstances of *appropriate* application’, as it needs to if Brandom is to avoid hyperinferentialism. It is not appropriate to apply the concept *red merely* because my eyes are open and it is daylight. There must also be a red thing in front of me!

inferences, which integrates them into the inferential role of at least those sentences that are capable of being noninferentially elicited by an RDRD.²⁰

By incorporating RDRDs into this theory Brandom thus aims to provide a more plausible *strong* inferentialist account of empirically-contentful language that gives a semantic role to the connection between language and the world, *without* having to appeal to representational relations such as reference or denotation. Thus, although Brandom is not interested in perception for its own sake—he boasts that the word ‘perception’ does not feature in the remarkably thorough index of *Making it Explicit*—perceptual circumstances of application play an important role in his semantic theory as the means by which inferential roles come to incorporate things outside the realm of language without wandering into the territory of representational semantics.

In the next section I argue that despite Brandom’s claims to the contrary, his account is actually *not* a version of strong inferentialism but rather of *hyperinferentialism*. I then propose an alternative view of the inferential significance of circumstances of application that, I claim, avoids this pitfall, and use this to motivate an account of perception, which I explore in the remainder of the chapter.

3.3 From Language to Perception

3.3.1 Brandomian Inferentialism is a Hyperinferentialism

We saw in the previous section that the difference between strong inferentialism and hyperinferentialism is that according to strong inferentialism the inferential role of a sen-

²⁰All that is required to avoid hyperinferentialism is that not *all* sentences have purely linguistic inferential roles. Some concepts and sentences, even ones with empirical content, may not have observational uses. Cf. Brandom, 2002a.

tence includes not just inferential connections to other sentences but also to *nonlinguistic* circumstances and consequences of application, such as the visible presence of a red thing or the action of eating an apple. The problem I want to raise for Brandom's account concerns the role that Brandom gives to RDRDs in his story about how he avoids hyperinferentialism.

We can begin to see the shape of the problem by considering *why* Brandom mentions RDRDs in the first place. RDRDs, after all, are purely causal mechanisms: Even thermometers and cats have reliable dispositions to respond differentially to changes in their environment. Reliable dispositions might play a role in a *descriptive* use theory of meaning such as Horwich's, which grounds its semantics in merely causal features of use such as the dispositions of speakers to accept sentences under certain conditions. But inferentialism is a *normative* use theory of meaning, so it is not immediately clear what role RDRDs have to play in an inferentialist account.

As I noted above, Brandom claims that circumstances of application have a role to play in *attributions* of reliability. But this doesn't give the circumstances themselves a normative or inferential role either. Consider an analogy: If Benedict's cat Pangur only meows when it is hungry then Benedict might treat his cat as a reliable responder to hunger by endorsing the inference from his cat's meowing to (the propriety of endorsing) the claim that Pangur is hungry. This clearly has no implications for whether there are inferential connections between *Pangur's hunger* and his claim that Pangur is hungry, or that the circumstances under which Pangur meows are part of the inferential role of the sentence "Pangur is hungry". It only shows that there's a good inference from the *claim* that Pangur meows in circumstances in which he is hungry to the claim that Pangur is hungry. But this much is true even on a hyperinferentialist account.

Brandom himself happily acknowledges that RDRDs themselves are merely causal mechanisms and so, *to that extent*, do not feature in the inferential roles of sentences. Their relevance to inferential semantics is derived from the normative significance of the *responses* reliably elicited:

“noninferential reporting practices involve two distinguishable components: a reliable differential responsive disposition and the capacity to exercise that disposition by endorsing a claim. *It is this second component that puts the response into the inferentially articulated space of applications of concepts, by bringing it into the game of giving and asking for reasons.*” (Brandom, 1994a, p.293, my emphasis)

Thus, it is only because a perceiver reliably responds to the presence of red things by *judging or asserting* that they are red that RDRDs play a *semantic* role in the inferentialist story. Pangur might respond to his hunger by uttering a cry, but this is of no semantic import because his cry cannot serve as a premise in an inference—it does not count as a *reason* for anything. But when a language using creature responds to the visible presence of a red thing by, say, uttering the sentence “That’s red”, this *can* serve as a premise in an inference, and so *is* normatively and semantically significant.

The problem with this response is that it is irrelevant to the question of whether Brandom’s view is a version of strong inferentialism or hyperinferentialism. What is needed for this theory to avoid hyperinferentialism is for the observational circumstances *themselves* to play an inferential role. But all that the invocation of RDRDs shows is that certain *causal responses* to these observational circumstances play an inferential role. As far as Brandom’s story is concerned, only the *responses* elicited by RDRDs—assertions, judgements, and the like—have any normative significance, or stand in any normative inferential relations to anything else. The observational circumstances themselves that elicit such responses, such as a red thing’s being visibly present to a perceiver, play no *normative* role whatsoever.

But if all of this is right—if on Brandom’s story observational circumstances have no normative significance, and so do not stand in any normative inferential relations to anything else—then we lose our right to say that the connections between circumstances and consequences of application are inferential *even in a broad sense*. Claims that are reliably elicited by a speaker in some circumstance, and even claims that *describe* the circumstances in which they were elicited, might bear inferential connections to language, but the actual observational circumstances themselves do not.²¹ Thus, Brandom’s semantic account is not, despite his claims to the contrary, a version of strong inferentialism at all but is, rather, a form of hyperinferentialism.

Brandom thus faces a dilemma: On the one hand, he *could* avoid hyperinferentialism by claiming that the role that exercising an RDRD gives to circumstances of application counts as instituting an inferential connection in a broad sense, but only by giving up on the idea that semantically-constitutive inferential connections are always *normative* connections. On the other hand, Brandom can hold on to the idea that semantically-constitutive inferential connections are always normative, but then in order to avoid hyperinferentialism he would have to give a genuinely *normative* role to circumstances of application. This is the option I pursue below.

The idea that Brandom misses *something* important by attempting to substitute RDRDs for perception proper is not in itself new (cf., e.g., McDowell, 2002). What is novel in the criticism I have provided in this section is that I have tried to argue that this oversight is of

²¹In *Between Saying and Doing* Brandom complicates his story by introducing the notion of a “Test-Operate-Test-Exit (TOTE) cycle of perception, performance, assessment of the results of the performance, and further performance—that is, a process or practice consisting of an open-ended sequence of feedback-governed performances” (Brandom, 2008, p. 178). I do not discuss this here because whilst TOTE cycles provide a greater degree of causal constraint than RDRDs whose operations are treated as distinct from one another, it does not solve the basic problem sketched here, since the inputs and outputs are still conceived purely causally. Creating normative links between causal inputs and causal outputs does not give a normative role to the world any more than the basic RDRD picture does.

significance not (just) for reasons internal to the philosophy of perception or epistemology, but rather for reasons relevant to the core tenets of the inferentialist view of language. Avoiding hyperinferentialism is a desideratum for inferentialist *semantics*, not (just) the philosophy of perception or epistemology. So if the criticism above is right, we need to find a replacement for RDRDs for purely semantic reasons, regardless of whether this account is satisfactory by other standards.²² In the next section I explain how I think we ought to respond to this problem.

3.3.2 Strong Inferentialism and the Space of Reasons

The basic idea driving the account of perception I provide in the rest of this chapter is suggested by Brandom when he says that what we need to acknowledge to avoid hyperinferentialism is that “the visible presence of red things *warrants* the applicability of the concept red” (Brandom, 2010, p. 658, my emphasis). I argued above that the normative aspect, indicated here by the use of the term ‘warrant’, gets lost in Brandom’s account as a result of the role he gives to RDRDs in his explanation. But the idea itself is a good one. What we need to acknowledge to avoid hyperinferentialism is that perception provides a *normative*, and not merely *causal*, constraint on the applicability of concepts—that perception is a *reason* for beliefs and not merely a *cause* of them. The contribution perception makes to semantics is not exhausted by the causal role it plays in eliciting assertions and generating beliefs, as it would be if it were no more than a particular kind of RDRD. Rather, it visibly presents us with things in such a way that it warrants the applicability of certain concepts and thereby supplies us with reasons for belief.

²²As Brandom notes (Brandom, 2002a, p. 99), McDowell’s criticisms are at least in part semantic, but McDowell does not attempt to spell the relevant semantic problem out as a tension *within* inferentialism, as I have in the current section.

Acknowledging a normative role for perception is sufficient to avoid hyperinferentialism, for it establishes genuinely *normative* connections between concepts and the circumstances that warrant their application. If a perceptual state in which one has something red visibly present to one *warrants* the applicability of the concept red—rather than merely *causing* the perceiver to apply the concept (though it may also do that too)—or *entitles* a perceiver to a claim—rather than merely *causing* the perceiver to form a belief—this means that there are *normative* relations between circumstances of application and concepts, not merely causal ones. And if there are normative relations between circumstances of application and concepts, we have a genuinely strong inferentialist semantics rather than a hyperinferentialist one.

In what follows, however, I want to draw out the implications of this idea not just for language but also for *perception* by providing an account of the content of perception that mirrors Brandom’s inferentialist account of metasemantics. The central feature of the inferentialist approach to metasemantics is a commitment to ***normative pragmatism***:

“Semantics must answer to pragmatics. The theoretical point of attributing semantic content to intentional states, attitudes, and performances is to determine the pragmatic significance of their occurrence in various contexts.” (Brandom, 1994a, p. 83)

Thus, when we say, for example, that an assertion has conceptual content, this is not *principally* to give it an empirical description—though it is to do that as well (after all, not just any old event counts as an assertion)—but rather to highlight the fact that assertions are the kind of thing that can be cited as reasons, and are the kind of thing for which reasons can be given. This is to give these events a certain normative status, indicating what role they should or may play in one’s reasoning, and in the (explicitly social) “game of giving and asking for reasons” (Brandom, 2000, p. 161):

“Claiming or asserting is what one must do in order to give a reason, and it is a speech act that reasons can be demanded for. Claims both serve as and stand in need of reasons or justifications. They have the contents they have in part in virtue of the role they play in a network of inferences.” (Brandom, 2000, pp. 161–162)

As the quotations above indicate, Brandom only applies this normative pragmatic approach to content to assertion and (intentional) action, but if we acknowledge that perception itself can warrant the application of concepts and serve as a reason for belief—as I have argued that we must if we are to avoid hyperinferentialism—then we can extend this approach to provide a foundational account of the content of *perceptual* states, too.

This is so far just a statement of purpose, however. To spell this idea out in more detail, we need to first provide a more thorough characterisation of how the normative role of assertions confers upon sentences their semantic content, and then show how this metasemantic account can be extended to provide a foundational account of the content of *perception*. I therefore turn to this task in the next section.

3.4 A Foundational Inferentialist Account of Perceptual Content

3.4.1 Inferentialist Metasemantics

The plan for providing the kind of account promised in the previous section begins with a more detailed explanation of the inferentialist *metasemantics* of natural language; how, or by virtue of what, it is that, on the inferentialist account, assertions and sentences come to have the content they do. By laying out this account in detail, we set up a framework for grounding attributions of content in normative pragmatic features of states and events that we can then apply to perception.

The normativity of the inferences that constitute the inferential role of sentences on the inferentialist account is grounded in the normative pragmatics that serves as its foundation. Specifically, the normative inferential role of a sentence is grounded in the way that assertions of that sentence alter the *normative statuses* of the speakers who assert them. There are two central statuses that form the core of the normative inferentialist account of pragmatics: *commitment* and *entitlement* (Brandom, 1994a, pp. 159–161). These are *deontic* statuses: To have a commitment is to have a certain kind of obligation or responsibility, whilst to have an entitlement is to have a certain kind of permission or authority.

For example, if Bruce signs a contract to play a gig, he thereby undertakes a commitment. This commitment confers on him a number of responsibilities: To play music at the venue for a certain amount of time, at a certain time, on a certain day, presumably of a certain quality. If Adam buys a ticket to see the show, he thereby acquires an entitlement. This entitlement grants him the authority to do things he would not ordinarily be able to do: He may enter the venue, stay there for (presumably at least) the duration of the event, use the facilities, and so on.

Bruce may be prevented from performing the gig, thus failing to honour the commitment he made by signing the contract, and Adam may be prevented from attending it, thus failing to avail himself of the entitlement he acquired by purchasing the ticket. But one does not, at least *ipso facto*, cease to be committed to something simply because one fails to honour the commitment. And one does not, again *ipso facto*, cease to be entitled to something simply because one fails to avail oneself of the entitlement.²³ Commitments spell out what

²³I include the *ipso facto* clause because obviously once the show is over Adam can no longer use the ticket to see the show. But this is not because he loses the entitlement that he had—as he might, for example, if he were banned from the venue after purchasing the ticket but before the date of the gig—but simply because it is no longer possible for him to do the thing the ticket entitled him to do. Some commitments are like this too—e.g. if I promise a friend to help him move to his new house, I will no longer be able to fulfil this commitment after the moving date has passed.

one *ought to* do, and entitlements what one *may* do, and this need not coincide with what one *in fact* does. Commitment and entitlement are thus essentially *normative* statuses.

As this example indicates, commitments and entitlements need not be linguistic. By asserting a sentence, however, a speaker commits herself *to a claim*. The normative significance of undertaking a specifically *linguistic* commitment of this kind consists in the distinctive way that assertions alter the normative status of the asserter. Specifically, assertions “involve alterations of deontic status *that have other alterations of deontic status as their consequences*” (Brandom, 1994a, p. 188, my emphasis). The normative statuses involved in the inferentialist account of the pragmatics of assertion are thus capable of serving as the foundation for an inferentialist semantics because they are, in a way to be spelt out in more detail below, *inferentially articulated*.

To illustrate, suppose that Benedict asserts the sentence “Pangur is a cat”, thereby undertaking commitment to the claim that Pangur is a cat. This event no doubt alters his normative status in a number of different ways. One aspect of the distinctively *linguistic* significance of this act concerns the commitments to *other* claims he thereby acquires. For example, although by uttering the sentence “Pangur is a cat”, Benedict *explicitly* undertakes commitment to the claim that Pangur is a cat, he also thereby *acquires* commitment to a host of other claims, including, e.g., the claim that Pangur is an animal, that there is at least one cat, and so on.

These normative facts about what *other* commitments a speaker acquires commitment to by virtue of asserting a given sentence therefore determine the *meaning* of that sentence, for the relations of inheritance that hold between these commitments suffice to determine the inferences that, on an inferentialist semantics, constitute that sentence’s particular inferential role. The inference from the claim that Pangur is a cat to the claim that Pangur is an

animal is a good one, and so is part of the inferential role of the sentence “Pangur is a cat”, *because* commitment to the claim that Pangur is a cat carries with it commitment to the claim that Pangur is an animal.²⁴ The former, semantic, fact is grounded in the latter, normative pragmatic, one.²⁵ In this way the inferentially articulated relations between normative statuses such as commitment suffice to ground the inferential role of sentences.

Thus, on this normative pragmatic inferentialist metasemantics, the non-semantic properties of assertions and the semantic properties of sentences are bound together by means of normative properties of agents. Specifically, the *semantic* properties of *sentences* are determined by the *normative* statuses of *speakers*, and these in turn supervene on the *non-semantic* properties of their *assertions*; namely, what words the speaker uttered. We can thus picture this part of the inferentialist metasemantics as in Figure 1.

At a reasonably high level of abstraction we can thus highlight two necessary conditions for an event or state to come to bear semantic content, on this inferentialist story: Firstly, the states or events must have the right kind of non-semantic properties. For example, sentences contain words such as ‘dog’ that can be distinguished from other words such as ‘cat’ so as to account for the fact that asserting the sentence ‘Pangur is a cat’ alters the speaker’s normative statuses in a different way than asserting the sentence ‘Pangur is a dog’.²⁶ Secondly, such states or events must affect the normative statuses of speakers in such

²⁴It is worth emphasising here that this is only *part* of the inferentialist account of pragmatics and metasemantics. But the part of the pragmatics given here is sufficient to ground the part of the semantics we are concerned with here; namely, the inferential role of a sentence as a *premise* in inferences. For the full story, including the explanation of the pragmatic basis of the inferential role of a sentence as a *conclusion* in inferences, see (Brandom, 1994a, ch.3)

²⁵We can now appreciate the full significance of Brandom’s claim that “what inferences are socially appropriate determines the meanings of the sentences involved in those inferences” (Brandom, 1983, p. 644), which I mentioned in §3.2.1 above.

²⁶*Exactly* what the ‘right kind’ of non-semantic properties are is an open question. As the sentence above indicates, they must be properties that have enough fineness of grain for speakers to be able to respond to differences in these properties by attributing different sets of commitments and entitlements. One might also suspect that they must exhibit certain combinatorial properties that explain how properties of parts come together to systematically determine the properties of wholes. Given that the problem of compositionality is a massive can of worms I am going to ignore this problem, and I am going to simply assume below that




Figure 1: The Normative Pragmatic Inferentialist Metasemantics

a way as to ground the attribution of *inferential* contents. Specifically, they must confer upon the speakers in question statuses like commitments that can bear inference-like relations to other such statuses, in the manner described above. It is only because commitment to one sentence carries with it commitment to others that these normative statuses are fit to play a role in grounding genuinely *semantic* properties of sentences, otherwise we might merely be describing a game of chess, or a set of rules of propriety.

Once we state things at this level of abstraction, however, it is clear that, at least in principle, states or events *other than assertions* might be fit to play this role, and hence that states or events *other than assertions* might bear conceptual content. As long as the states or events in question have the right kind of non-semantic properties, and confer upon agents normative statuses of the right kind, we can apply the same kind of metasemantic story we told in this section to give a foundational semantic account for such states or events. The suggestion I want to develop in the following section is that *perception* can play a role analogous to assertions in a generalised version of the metasemantic story told above.

3.4.2 A Foundational Account of Perceptual Content

In the discussion of pragmatics in the previous section I focused exclusively on the way that by undertaking an assertoric commitment a speaker acquires further *commitments*; commitments that stand in relations of consequence to the claim she explicitly undertook commitment to. But as we noted commitment is only one of the two central normative statuses that form the basis of this normative pragmatics. Assertions also alter the *entitlements* of speakers. Whilst asserting a sentence doesn't *in itself* entitle a speaker to the claim she makes, *if* the speaker is entitled to that claim then this entitlement will be inherited

experiential properties are the right kinds of non-semantic properties. I address the issue how perceivers recognise the content of their experience on the basis of such visual features in §4.4.2.

by the other commitments she acquires by undertaking commitment to it (Brandom, 1994a, p. 179). For example, if Benedict asserts that Pangur is a cat, and he is entitled to this claim, then he will also be entitled to the claims that follow from it, such as the claim that Pangur is an animal.

This point bears on the dimension of *responsibility* that, in the discussion in §3.4.1, I noted is characteristic of commitments in general. Specifically, by undertaking an assertoric commitment speakers acquire the *conditional task responsibility* of demonstrating their entitlement to their claim that they asserted, should their entitlement to this claim come into question (Brandom, 1994a, p. 172). Speakers can demonstrate their entitlement in a number of ways, the most obvious of which is by offering *reasons* for their belief; asserting other claims their entitlement to which entitles them to the claim in question. A consequence of the fact that entitlement to the claim that Pangur is a cat entitles Benedict to the claim that Pangur is an animal is that Benedict can offer the former claim as a *reason* for the latter, if challenged to demonstrate his entitlement.

On pain of regress, however, not *all* entitlement can be inherited (Brandom, 1994a, p. 204). There need to be at least some entitlements that speakers have that they have not inherited from some other source.²⁷ Brandom's solution to the problem posed by the threat of regress is to appeal to the notion of *default entitlement*. In the case of perception, specifically, he claims that speakers have a default entitlement to observation reports; assertions that are elicited through the exercise of an RDRD (Brandom, 1994a, p. 222). This allows observation reports to serve as regress stoppers because the default entitlement they have can only be overturned by means of a *reasonable* challenge—that is, someone's entitlement to a claim

²⁷This applies even if we consider sources of entitlement other than explicit justifications, such as testimony. See Brandom, 1994a, p. 204.

can only be overturned by a challenge to which the challenger is herself entitled (Brandom, 1994a, p. 178).

Brandom's solution to this problem is thus coloured by his commitment to reduce the contribution of perception to that of a mere reliable belief-forming mechanism. One consequence of this fact is that Brandom focuses almost exclusively on questions of entitlement to commitments *that speakers have explicitly undertaken*, in this case in the form of an observation report. A perceiver exercises an RDRD by issuing an observation report, and only at *that* point does the question of their entitlement to that claim arise.

But commitment and entitlement are statuses of *people*, not states or events. And outside of language we often encounter situations in which people are entitled to do things that they have not done yet, or will not do ever (cf. Brandom, 1994a, pp. 161–2). If someone buys a ticket for a particular train journey they have acquired the entitlement to take that train, whether they actually do or not. Once somebody actually gets on a train the question then arises of whether they were entitled to do so or not, but the person who purchased their ticket ahead of time had their entitlement before they ever actually availed themselves of it by stepping onto the train.

On Brandom's account there is no parallel to this sort of situation in the case of language. Only claims speakers have *already* undertaken commitment to have the kind of default entitlement characteristic of observation reports. But it is the *person* who is entitled to say what she did, so why should we think that the question of her entitlement only arises once she has undertaken a commitment? The idea that a person might be entitled to claims that she hasn't yet made, or that one speaker might *attribute* entitlement to a claim to another speaker without *also* attributing commitment to that claim, is simply absent from Brandom's account. Brandom thus treats entitlements as just a privileged kind of commitment:

“[P]layers of the game of giving and asking for reasons must also distinguish, *among the commitments an interlocutor undertakes* a distinguished sub-class to which she is *entitled*.” (Brandom, 2008, pp. 113–114, partly my emphasis)

Were it not for Brandom’s commitment to think of perception as merely a reliable mechanism for producing observation reports, I think a much more natural answer to the regress problem would suggest itself: It is *perception* that serves as the regress-stopper, by providing a perceiver with entitlements; entitlements that she may, but need not, choose to avail herself of. If she *does* avail herself of this entitlement, by making a claim, the question then of course arises whether she *was* so entitled. But, like the person who buys a train ticket before she takes her journey, she was entitled to this claim even before she actually asserted it, and would have been entitled to do this even if she never availed herself of this entitlement.

If we accept the idea, suggested in §3.3.2, that experiences are sensory states—states of having things visibly present to one—that warrant the application of certain concepts, we can accept this much more natural solution to the regress problem. When a perceiver applies the concept apple in response to the visible presence of an apple—say, by asserting the sentence “There’s an apple on the table”—she does not acquire entitlement to this claim *by* exercising a reliable disposition to produce observation reports in response to apples, but rather merely *avails herself* of an entitlement that her experience itself already conferred upon her.²⁸

As I am using these phrases here, to be warranted in applying a concept to a given object is to be entitled to believe or assert a claim that applies that concept to the object

²⁸A commitment that a perceiver undertakes on the basis of her experience might fail to inherit this entitlement if the perceiver does not form the relevant belief in the right kind of way. But this is generally true of inheritance of entitlement: Benedict’s entitlement to the claim that Pangur is a cat also entitles him to the claim that Pangur is an animal, but a commitment only inherits this entitlement if it is formed on the basis of this former belief, and so can fail to be inherited if Benedict does not form his belief in the right way. This relates to the point discussed in §2.4.1 that to characterise a person’s perspective we need to acknowledge not just normative relations but specific instances of these relations.

in some manner. As I mentioned above, application can occur in predication, as when a speaker attributes the concept red to an object by saying “That is red”. But application can also occur in noun-phrases, as when a speaker applies the concept red to an apple by saying “That red fruit is an apple”, and in quantified expressions, as in “Benedict ate the red apple” (cf. Burge, 2010, pp. 31–32).

The particular claims a perceiver is entitled to will correspond to the particular constellation of concepts her experience warrants the application of, and their relation to one another. Ordinarily experience will warrant the application of a number of concepts with respect to the same object, so that, e.g., the visible presence of a red apple will warrant the application of *both* the concept red *and* the concept apple, thus entitling the perceiver to the claim that that apple is red, that that red thing is an apple, and so on for other possible claims in which those concepts are applied. In the rest of this chapter, therefore, I will continue to use these two vocabularies interchangeably, on the assumption that they are expressively equivalent.²⁹

These observations have implications not only for thinking about assertion but also for thinking about perception. Given that experiences, like assertions, alter the normative statuses of perceivers that undergo them, it is possible for us to tell the same kind of metasemantic story about experiences as we told about assertions in the previous section. Perceptual experiences have semantic contents for the same reason that assertions do; namely, because the normative statuses a perceiver acquires by undergoing an experience are

²⁹This explanation is offered primarily as a way of squaring this talk of application of concepts with the theoretical priority of propositions on the inferentialist account. One consequence of this point, however, for an account stated purely in terms of entitlement to claims is that the issue of entitlement need not be all-or-nothing. A perceiver might see a small part of what is in fact a red apple, but not be in a position to make out what kind of fruit it is. If someone asks her what’s on the table and she takes a guess and says “That’s a red apple”, her experience partly entitles her to this claim, in that it warrants the applicability of the concept red, but it doesn’t entitle her to the claim *tout court*, since it doesn’t warrant her application of the concept apple, given her limited view of the object. This relates to the issue of the inferential dimension of the intentionality of perception, which I discuss briefly in §3.5.2 below.

inferentially articulated, and so are fit to confer upon her experience inferentially articulated content. The difference is merely that assertions confer *commitments* on asserters, whereas perception confers *entitlements* on perceivers.

In a paradigmatic case, a perceiver is visually aware of an object from a certain distance, at a certain angle, in certain lighting, and so on. By virtue of being in a perceptual state of this kind, the perceiver thereby has *entitlement* to various claims. For example, if a speaker is visually aware of a red apple on a table (in good viewing conditions), this experience may confer on her entitlement to the claim that there is a red apple on the table, *and by virtue of conferring this entitlement on the perceiver she will acquire* entitlement to various other claims including, e.g., the claim that the apple is coloured, that there is some fruit on the table, and so on.

These normative facts about the claims the perceiver is entitled to by virtue of being in that perceptual state thereby determine the *content* of her experience, for the relations of inheritance of entitlement that hold between these entitlements suffice to determine the inferences that, on an inferentialist semantics, constitute that experience's particular inferential role. That an experience in which a perceiver is entitled to the claim that there is a red apple on the table also confers on the perceiver entitlement to the claim that there is a fruit on the table thus partly determines the inferential role of that experience. The inferential role of the experience is grounded in the way it changes the normative status of the perceiver, just as the inferential role of an assertion is grounded in the way *it* changes the normative status of the perceiver.

Thus, by direct analogy with the case of assertion, the non-semantic properties of perceptual states and their semantic properties are bound together by means of the normative properties of perceivers that undergo them. Hence the *semantic* properties of *experiences*

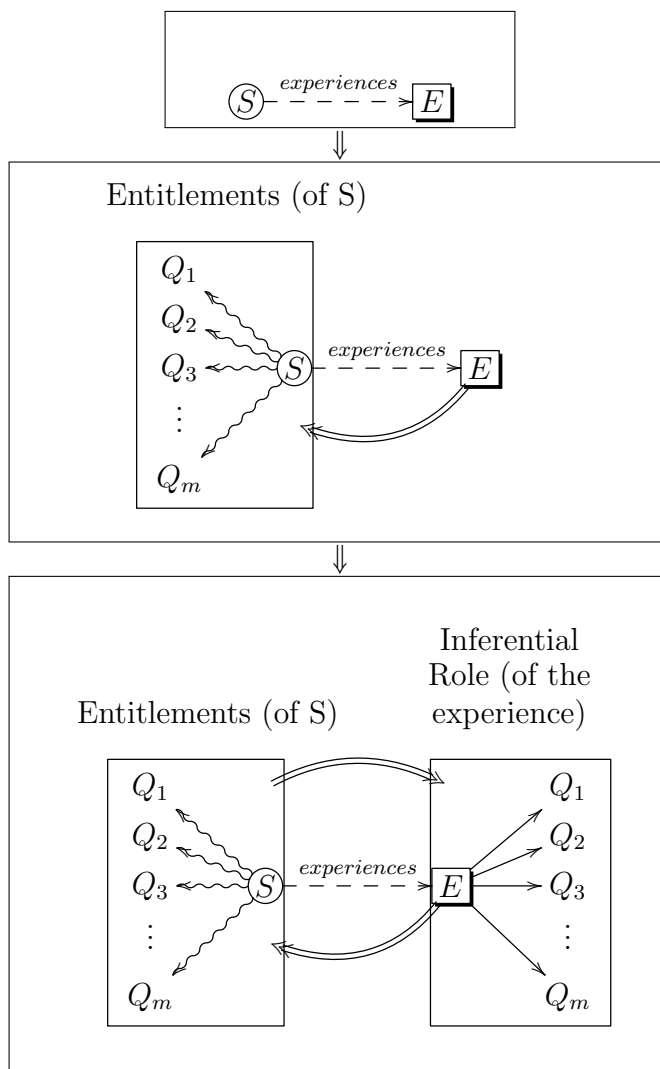


Figure 2: The Normative Pragmatic Account of Perceptual Content

are determined by the *normative* properties of *perceivers*, and these in turn supervene on the *non-semantic* properties of experiences; namely, what the perceiver was visually aware of. And so by analogy with assertion we can picture this relationship as depicted in Figure 2.

In this way, the inferentialist metasemantic story developed above can be extended to provide a foundational semantics for *experiences* in addition to sentences. Just as the content of a *sentence* is constituted by its inferential role, so too the content of an *experience* is constituted by *its* inferential role. That perceptual experiences have (inferentially articulable) content is thus a consequence of the fact that they alter the normative statuses of perceivers, given that, on the normative pragmatic account of content, the conceptual content of states is constituted by their normative inferential role. By giving perceptual circumstances of application a role in constituting the inferential content of sentences, we can thus not only give a more satisfactory semantics for *sentences*, as I argued in §3.3.2, but we provide a framework that can be applied to explain the content of *perception*, too.

Before I end this section, however, I want to highlight two features of this account that bear on how we think of the epistemology and metaphysics of perception—namely, my choice of *entitlement* as the relevant normative status, and my description of perception as a *state* rather than as an event, like assertion.

Firstly, I want to acknowledge that perception may *also* confer commitments on perceivers that they have not explicitly *undertaken* commitment to by judging the relevant thing. Arguably, if a perceiver has a red object clearly in her line of sight, and she has no reason to believe that the lighting conditions are unusual, or that her perceptual capacities are malfunctioning somehow, she is not merely *entitled* to judge that the object is red, she is *rationally compelled* to judge that it is red (Gupta, 2019, §79). And this is plausibly so even

if she does not actually *judge* that it is red. Thus, the description of the normative significance of perception I have given here does not aim to provide a *complete* characterisation of the normative significance of perception. It is enough for this account to succeed that perception plays *at least* the normative role described above, even if it also plays other roles.

The reason why I use entitlement rather than commitment or some other kind of agentive normative status (cf. Simonelli, forthcoming) is because entitlement is the *broadest* kind of normative status that perception confers upon a perceiver. If the content of perception were determined only by the claims a perceiver is *obligated*, rather than *permitted*, to make then huge amounts of the detail of visual experience would not be represented in the content of perception. For example, if I am aware of a red book I might be compelled to judge that it is red, but I am not compelled to judge that it is the specific *shade* of red it is. If the content of perception is determined by the claims a perceiver is *entitled* to make, by contrast, then we get a much richer and more plausible account of the content of perception, for plausibly my experience *entitles* me to claims about its specific shade, even if it does not *commit* me to such claims.

This brings us to our second point; the characterisation of perception as a state rather than as an event. Although states are usually treated as a sub-class of events, the reason why I described perception as a state rather than an event *tout court* is that, in the context of an analogy with assertion, calling perception an event would be misleading. Assertions are, plausibly, *achievements* (cf. Ryle, 1949, pp. 131–135): They are, logically speaking, instantaneous events. It makes sense to ask “For how long did he speak?”, but not “For how long did he *assert*?”. By contrast, perception is, most plausibly, a *state* (cf. Vendler, 1957, pp. 146–7): It is a homogeneous event that extends over time and does not culminate in an achievement. By contrast with assertions, it *does* make sense to ask “For how long did he

perceive the object?” or “For how long was he perceptually aware of the object?”—though not “For how long did he *see* the object?” (Ryle, 1949, pp. 134–135; cf. Vendler, 1957, p. 147).³⁰

The reason that this point is important is because it is only by acknowledging the temporally extended nature of perception that we can make sense of the idea that perception can entitle perceivers to claims that they never undertake commitment to. For example, suppose that a perceiver sees a scarlet object and judges that it is red but does not judge that it is scarlet. As long as she is still perceiving the object she will be entitled to the claim that it is scarlet, but once she looks away she will lose her entitlement to this claim.³¹ We could not accommodate this kind of situation if perception were an achievement, for we would either have to say that the perceiver acquired the entitlement *tout court*—which is unacceptable given that, unless she undertakes commitment to the claim, she loses her entitlement once she stops perceiving—or that she did not *tout court*—which is unacceptable because it fails to capture the fact that the perceiver *did* have entitlement to the claim *before* she stopped perceiving.

Thus, although I will sometimes speak of perception as an event, it is important to keep in mind that it is a *temporally extended* event; specifically, the *state* of being visually aware of something.

As these two discussions demonstrate, although I have principally motivated this foundational account of content by means of *semantic* considerations, it bears in important ways on issues in the metaphysics, epistemology, and phenomenology of perception. Some issues,

³⁰One might *observe* an object or *watch* an event unfold for some period of time, but not *see* it for a length of time. We should thus be careful to distinguish observing from seeing, even if visually observing always entails seeing (cf. Ryle, 1949, ch. 7). The same applies, *mutatis mutandis*, for perceiving.

³¹It might be possible to imaginatively recall the experience for a short period of time after the experience, and if it is possible to imaginatively recall an experience in this way without having formed a belief about what one saw then one’s entitlement might persist for some time after the experience has ended. Whether or not this is possible the current point still stands.

such as the two just discussed, *need* to be resolved in certain ways if the account is to succeed at all, or at least to be sufficiently plausible. However, this is not true of the majority of questions in these areas.

In fact, it is perhaps even helpful to think of this account as giving a *framework* for foundational semantics, for the explanation I gave is compatible with a number of different specific views of the metaphysics, phenomenology, and epistemology of perception. In part this is an obvious benefit, making the framework more widely applicable than it would be if it depended on specific and possibly tendentious answers to questions in these areas. One implication of this feature of the account, however, is that different answers to questions in these other areas could lead to different views of exactly what the conceptual content of perception is, and how it relates to the metaphysics, phenomenology, and epistemology of perception, some of which might be more plausible than others.

In this respect the explanation given above is similar to other frameworks for foundational semantics such as teleosemantics. Teleosemantic views all share certain features—namely, they all give an explanation of the content of perception that grounds the accuracy conditions of perceptual states in the proper function of the systems that produce or consume them. But there are different ways of cashing out this general approach, depending on whether one focuses on the proper function of systems that *produce* perceptual states (Matthen, 1988) or systems that *consume* them (Millikan, 1989), whether the view bases its account on an informational conception of perceptual states or not (Dretske, 1995), and how one explains the notion of proper function and its relation to biological function (Dretske, 1986). These different versions of teleosemantics often deliver different verdicts about what a creature represents in specific cases, and sometimes disagree about what kinds of properties can be represented in perception in general, but they all count as versions of teleosemantics because

the explanations they give are all in some way rooted in teleological notions such as proper function.

In the next chapter I take a stance on some important issues regarding the metaphysics and phenomenology of perception, but in these final few sections of this chapter I want to suggest one addition with which one might fill out the framework I gave above that does not depend on the views developed in the next chapter. Specifically, I address the question of how to think about the conceptual, inferential significance of awareness of *objects* on this account.

3.5 Givenness and the Inferentialist View of Perception

3.5.1 Object-Awareness and The Myth of the Given

Above I have argued that being visibly presented with an object can warrant a perceiver in applying concepts to that object, and that this fact can ground an explanation of the content of her experience. One problem for this kind of explanation is that it seems to fall into what Sellars calls the Myth of the Categorial Given. Here is how Sellars describes this Myth in the *Carus Lectures*:³²

“If a person is directly aware of an item which has categorial status C, then the person is aware of it *as* having categorial status C. This principle is, perhaps, the most basic form of what I have castigated as ‘The Myth of the Given’.” (Sellars, 1981, §44, p.236)

³²The dialectical context within which Sellars spells out the Categorial Myth is importantly different from the current one, and actually occurs as part of Sellars’s attempt to argue, against Roderick Firth, that concepts of visible features of objects are conceptually more basic than concepts of features of experiences (cf. Sellars, 1981, §§38–76, pp. 235–243). Nonetheless I think that, as is often the case with Sellars, his general point applies much more broadly. The solution I give is also in many respects different from the account Sellars gives in the *Carus lectures*, but it bears an important similarity to McDowell’s Sellarsian view developed in his *Woodbridge Lectures* McDowell, 1998a, 1998b, 1998c, as I note below.

“To reject the Myth of the Given is to reject the idea that the categorial structure of the world—if it has a categorial structure—imposes itself on the mind as a seal imposes an image on melted wax.” (Sellars, 1981, §45, p. 237)

To fully address concerns about the Myth of the Given will require addressing the question of how it is that perceivers come to be capable of having inferentially-articulated entitlements *at all*—a question that I address in §3.5.3. However, here I want to focus on the specific question of what it means for perception to present a perceiver with an object *as* an object. The account I have given above seems to depend on the idea that merely being perceptually aware of what is in fact a red apple is sufficient to entitle a perceiver to claims such as the claim that there is a red apple on the table. Thus, it seems as if I have suggested that by being perceptually aware of a mind-independent object, a perceiver is thereby guaranteed to be aware of it *as* a mind-independent object which, if true, would entail that the account falls straightforwardly into the Categorial Myth of the Given.

Even if we were not concerned by the Categorial Myth, however, it would be an unhappy feature of the account if it assumed that being perceptually aware of an object entails being entitled to attribute concepts to that object *as* an object. Hannah Ginsborg, for example, gives an example that seems to show that being presented with an object and some of its visible features is not sufficient to have reason to believe *that* the object has those features, and hence to see it *as* an object (cf. Gupta, 2019, § 54; Siegel, 2010, p. 47):

“On the face of it... your merely seeing the package does not seem to be the kind of thing which could rationalize [that is: give you a reason for] the belief in question. For suppose that, on seeing the package, you take it to be, not a package, but rather a patch of sunlight or a pile of newspapers; or that you merely register its presence without taking it to be anything in particular at all. Even if your perception causes you to form the belief that your books have arrived, say by an unconscious process of association, there does not seem to be any ground for supposing that the process of belief formation counts as rational in the light of your perception.” (Ginsborg, 2011, p. 136)

This was something that I glossed over in characterising the relationship between being warranted in applying a concept to an object and being entitled to certain claims. In §3.4.2 I suggested that the vocabulary of being warranted in applying a concept and having entitlement to claims can be used interchangeably. But Ginsborg’s point threatens to sever this connection, or at least make it much less clear how the two notions relate. I might be visibly presented with the shape of the package, and so warranted in applying the concept square, say, and visibly presented with its colour, and so warranted in applying the relevant colour concept, but *not* for that reason alone entitled to the claim that the *package* has that colour. How can we explain this, on the inferentialist account?

Ordinarily, views that attribute content to perception think of the worldly entities of which a perceiver is aware as the *intentional objects* of perception; as “what the intentional act is about or directed at” (Crane, 2006, p. 207). When Benedict thinks that Pangur is a white cat, *Pangur* is the intentional object of his thought, and he thinks *of* Pangur *that he is a white cat*, which is the content of his thought. So to understand how the inferentialist account ought to deal with this kind of case, and how it avoids the Myth of the Categorical Given, we need to understand how to think about intentionality on the inferentialist account.

3.5.2 Intentionality and Inference

Brandom’s account of intentionality is incredibly detailed, and applying it in full to perception would take more space than we have available here, but we can at least outline enough of the view to solve the problem raised in the previous section. Brandom’s strategy is to answer the question of what it is to recognise an object *as* an object by answering the question of what it is to use an expression as a singular term; that is, by explaining what

purporting to *refer* to an object amounts to. Brandom, following Frege, suggests that the answer to this question involves appeal to the notion of criteria of identity:

“That a judgment is directed toward an object is intelligible only in the context of practices of identifying objects as the same again, and individuating them as distinct... This is to say that the use of expressions as singular terms essentially involves, not only norms that could be made explicit as criteria of *application*, but also norms that could be made explicit as criteria of *identity*.” (Brandom, 1994a, p. 416)

Perceiving an object, I have argued, requires that a perceiver grasps, at least implicitly, norms that could be made explicit as criteria of application, as one might make explicit the norms governing the application of the concept red by stating a rule of application such as “Only apply the concept red (noninferentially) to visibly red things”. But Brandom suggests here that the intentionality of judgement—its directness at an object—involves a grasp not of norms that could be made explicit as criteria of application but norms that could be made explicit as criteria of *identity*.

Norms that could be made explicit as criteria of application are norms governing the application of concepts; the conditions under which one is permitted or obligated to apply a concept. Norms that could be made explicit as criteria of *identity*, by contrast, are norms governing the *intersubstitutability of singular terms*; the conditions under which one is permitted or obligated to substitute one singular term for another. Like the norms governing the application of concepts, these norms are to be understood in terms of their inferential role (Brandom, 1994a, p. 415).

Substitution inferences are a specific class of inferences in which the conclusion is obtained from the premise by substituting one term for another (Brandom, 1994a, p. 370), as in the inference from “George Orwell wrote *1984*” to “Eric Blair wrote *1984*”. What is semantically distinctive about singular terms is that the substitution inferences they feature

in are *symmetric*: If the aforementioned inference is good then *ipso facto* so is the inference from “Eric Blair wrote *1984*” to “George Orwell wrote *1984*” (Brandom, 1994a, pp. 371–372). To treat these two sentences as referring to the same person thus *just is* to take oneself to be warranted in substituting the term “George Orwell” for “Eric Blair”, and vice versa. To do this is to implicitly grasp norms that could be made explicit by means of an *identity* claim such as “Eric Blair *is* George Orwell”. If one takes two terms to be substitutable in this way I will say that this involves *commitment to (symmetric) substitution inferences involving those singular terms*.³³ (The same applies, *mutatis mutandis*, to claims and singular concepts rather than sentences and singular terms.)

Although this discussion so far merely concerns the issue of what kinds of inferences one has to grasp to be said to *treat* two terms as referring to the same object, it carries implications for the issue of what kind of normative significance perception must have if it is to be capable of presenting a perceiver with an object *as* an object. If to treat two claims as being about the same object is to be committed to substitution inferences involving the singular concepts in those claims, then, intuitively, for experience to have as part of its content two claims that are *about* the same object is for it to *entitle* the perceiver to substitution inferences involving the singular concepts in those claims. And so to be presented with an object *as* an object at all is to be entitled to at least *some* substitution inferences.³⁴

This characterisation of the intentional aspect of perceptual content is admittedly fairly abstract, but to illustrate how this suggestion can help explain what is required for an experience to present an object as an object, and hence how an experience can *fail* to

³³This should not be taken to imply that we could think of these terms *as* singular terms independently of their involvement in symmetric substitution inferences. It is perhaps more natural to say that it involves commitment to the *identity claim* rather than commitment to the *inferences*, but I do not want to presuppose that the perceiver explicitly grasps the concept of identity, for reasons I discuss below.

³⁴These inferences might be as basic as the inference from the claim that *that thing I saw a second ago was yellow* to the claim that *that thing I am seeing right now is yellow*, or from the claim *that patch of light is yellow* to the claim that *that patch of light is bright*.

present an object as an object, let us use a more detailed version of an example similar to Ginsborg's: Consider Dick, who is looking at the base of a tree where there lies a hoodie on which is printed a camouflage design of leaves, twigs, and patches of various earthy colours. Given the realism of the design on the hoodie, and the dim light of the forest, Dick does not realise that he is looking at a hoodie, and in fact his experience does not present this to him *as* an (individual) object at all, but rather as a pile of mulched leaves, twigs, and soil of the kind that usually occupies forest floors.

Dick, by virtue of being aware of the various properties of the hoodie, is warranted in applying various concepts, such as certain basic colour and shape concepts, to the thing—whatever kind of thing it happens to be—that occupies that part of his visual field. He might articulate the content of some of the claims to which this experience entitles him by saying “*This is green*” or “*That is lumpy*”, where “this” and “that” refer to different parts of the hoodie.³⁵

Given that he does not see the hoodie *as* a hoodie, or even see the brownness and lumpiness as features of one object, Dick's experience does not entitle him to substitute one of these demonstratives for another. That is, his entitlement to the claim that *This is green* does not entitle him to infer that *That is green*, and his entitlement to the claim that *That is lumpy* does not entitle him to infer that *This is lumpy*. His experience warrants him in applying these concepts to *something*, but he is not warranted in applying them *together* in the way that would constitute awareness of the hoodie *as* a single object.

Suppose however that, after looking at the base of the tree for a while, Dick suddenly has a Gestalt shift, and he realises that he is looking directly at his missing hoodie. We can characterise this change as a change in the *conceptual content* of his experience by saying

³⁵The talk of reference here is not presupposing anything substantial. We can imagine that Dick is shifting his gaze from one part of the hoodie to the other as he utters these two sentences.

that his experience now not only warrants him in applying these concepts to *something*, but warrants him in applying these concepts to *the same thing*. In explicitly *inferential* terms, we can say that his experience not only entitles him to the claim that *This is green* and *That is lumpy*, but also entitles him to substitute each of these demonstratives for the other.

To undertake commitment to the content of this experience would thus be to undertake commitment not just to these two demonstrative claims but also to these substitution inferences; that is, to treat these terms as intersubstitutable and, hence, to treat the two claims as *about* the same object. If one explicitly grasps the concept of identity one could express this by asserting an identity claim, but one could also express one's awareness of it as an object by referring to it using an expression that presupposes the coreference of the demonstratives that feature in the claims to which one's experience entitles one, as Dick might by saying "That green lumpy thing at the base of the tree is my hoodie!".

The suggestion, then, is that we can make the notion of awareness of an object *as* an object conceptually respectable—fit to be part of an inferentialist account of the content of perception—rather than needing to think of it as some mysterious extra-conceptual feature of perception. Given merely by the fact that one is visually aware of an object, by thinking of perception as not merely entitling perceivers to various claims, but also as entitling them to *substitution inferences* involving the singular concepts in those claims.

The substitution inferences to which an experience entitles a perceiver capture the semantic connections between different parts of the visual field of the perceiver, grouping the claims to which perception entitles a perceiver into clusters bound together by symmetric relations of intersubstitutability—as claims, that is, that are all about *the same* object. This gives us a way of making sense, in distinctly inferentialist terms, of the kind of rational unity

of an experience of an object (*as* an object) that John McDowell describes, in Kantian terms, as *intuitional*:

“If an object is present to one through the presence to one of some of its properties, in an intuition in which concepts of those properties exemplify a unity that constitutes the content of a formal concept of an object, one is thereby entitled to judge that one is confronted by an object with those properties.” (McDowell, 2008a, p. 12)

On McDowell’s intuitional view perceptual experiences have content that can be expressed most perspicuously by means of a noun-phrase such as “That red apple there” (cf. McDowell, 1998a), and he contrasts this with the view he endorsed in *Mind and World* (1996), according to which experience has *propositional* content, content that can only be expressed fully by means of a *sentence* such as “That apple there is red” (McDowell, 2008a, pp. 6–7). But in light of the explanation I have given above, we can accommodate the idea that perception of objects *as* objects has a distinctive formal unity without having to give up the idea that the content of perception is propositional.³⁶

If perception of an object entitles a perceiver not just to claims but also to substitution inferences between those claims, we can explain the unity perception of an object *as* an object has without having to acknowledge anything other than fully propositional contents. As I initially suggested above, the vocabulary of being warranted in applying a concept—or, now,

³⁶It is worth noting one point of interest regarding the formal nature of this concept of an object. If what I have argued here is correct then in order to be able to perceive objects *as* objects one has to be able to deploy substitution inferences of this kind. But as Brandom notes (Brandom, 1994a, pp. 416–419), identity claims *explicate* substitution licenses, in the sense that if one can deploy substitution inferences one can do everything one needs to do to be able to make identity claims. Furthermore, Brandom and Sellars have an account of categories as concepts that explicate necessary features of language use (Brandom, 2015, ch. 1; Sellars, 1970). So, on the plausible assumption that one must be able to see an object *as* an object in order to be able to have (objective) experiences, identity is a category, in the Brandom-Sellars sense, that is required in order to perceive objects *as* objects—if you can *do* what it is necessary to be able to do in order to have objectively contentful experiences then you can already do everything you need to do to grasp the concept of identity. The connection with McDowell’s Kantian conception of intuitions is not merely a cute analogy, therefore, but hints at something deeper. This obviously only gestures towards the details of this idea, but I intend to explore it further and extend this basic approach to other Kantian categories in future work.

being warranted in substituting one singular term for another—is expressively equivalent to that of being entitled to claims—or, now, being entitled to substitution inferences.

Given that this is a purely *formal* conception of an object it is perfectly possible that a perceiver can be visually related to what is *in fact* a physical object without thereby seeing it *as* an object in this formal sense. Dick, prior to his Gestalt shift, is aware of a coat but not aware of it *as* a coat, nor even as a single object at all. And if the way things are arranged in one’s environment is suitably misleading, as sometimes happens with certain visual illusions, one can be aware of something that is *not* an object (or at least not a *single* object) *as* an object. Thus, this explanation allows us to avoid the Categorical Myth of the Given by distinguishing between being aware of what is in fact a physical object and being aware of it *as* an object.³⁷

3.5.3 Perception and Objects in Language Acquisition

In the explanation of how to avoid the Categorical Myth I gave in the previous section I took for granted the fact that a perceiver has an implicit grasp of the norms governing the application of concepts and norms governing the intersubstitutability of singular terms; that

³⁷One consequence of the formal nature of this conception of an object is that the part of the account of perceptual intentionality I have given here doesn’t tell us *when* a perceiver is aware of an object as a single object. If we stick to the orthodox inferentialist explanation of intentionality, however, we can accommodate this fact, for the question of what objects a perceiver is *actually* aware of will not be settled by what she *takes* herself to be entitled to, nor by some ‘objective’ facts about what experiences of that kind entitle perceivers to, but will instead be cashed out in terms of the *social* dimension of the normative statuses acquired in perception. Locutions such as “saw what was *in fact* a hoodie” play a role analogous to *de re* specifications of the content of assertions; they serve to express the intentionality of perception. If Harry sees Dick’s hoodie *as* a hoodie, and recognises that Dick *doesn’t* see it as a hoodie, he can express this by means of a locution such as “Dick sees what is *in fact* a hoodie *as* a pile of dirt and leaves”. Addressing this point in detail would take us too far beyond current concerns, but it is important to bear this in mind in the next chapter, where for dialectical purposes I simply assume, with my naïve realist interlocutors, that perception is directed at objects, for this kind of approach would allow us to respond to criticisms of naïve realism based on similar problems, such as the example of Gupta’s that I cited above in connection with Ginsborg’s case of the yellow box (Gupta, 2019, § 54). For Brandom’s account of the expressive function of *de re* locutions and their role in securing the objectivity of linguistic intentionality, see Brandom, 1994a, ch. 8.

their experiences warrant them in applying concepts and in substituting certain singular terms for one another. In this section I want to provide an explanation of how it is that I think perceivers come to have such entitlements, by virtue of the role that perception plays in the acquisition of language. In the next chapter I combine the normative pragmatic approach to content with an approach to the metaphysics and phenomenology of perception that takes *relations to mind-independent entities* to be central to perception. The discussion I give here of language-learning will also serve to motivate this account.

If perception supplies a perceiver with entitlements to claims or with warrants to apply concepts then that perceiver must be capable of *bearing* such entitlements, and in order to *bear* such entitlements she must be capable of *availing* herself of such entitlements. For perception to have content, then, given the way that this content is grounded in the way it alters the normative status of perceivers, it is not sufficient that the perceiver's sense organs merely be stimulated in a certain way, or that her synapses fire in a certain pattern. She must be a *language user*, and she must be capable of telling what sorts of linguistic entitlements her experience confers upon her.

Acquiring a language thus involves learning not just what *words* mean but, so to speak, what *experiences* mean; being taught the normative significance of both assertions *and* experiences. Here is how Sellars describes the importance of acquiring a language for linguistically-informed rational perception:

“coming to see something as red is the culmination of a complicated process which is the slow building up of a multi-dimensional pattern of linguistic responses (by verbal expressions to things, by verbal expressions to verbal expressions, by meta-linguistic expressions to object-language expressions, etc.) the fruition of which as conceptual occurs when all these dimensions come into play in such direct perceptions as that this physical object (not that one) over here (not over there) is (rather than was) red (not orange, yellow, etc.).” (Sellars, 1963, p. 334)

The process of initiating a child into a linguistic community thus involves not just teaching her how to respond appropriately to utterances—how to recognise what people are *saying*—but also teaching her how to respond appropriately to *experiences*—how to recognise what she is *seeing* (cf. Sellars, 1963, p. 334). At a minimum, this will involve language *users* inculcating in language *learners* patterns of linguistic responsiveness that they, the language users, know to be appropriate—e.g. by training children to respond to red things only by uttering (sentences that attribute the term) “red” (to that object) (cf. Sellars, 1969, 1974).

Although the very early stages of this process might involve a kind of brute reinforcement, when teaching a child a language—as opposed to teaching a dog how to sit, say—teachers will typically inculcate the relevant dispositions in the child by *treating it as if it can speak*, not by means of crude non-linguistic rewards and deterrents. (Though in cases of exhaustion on the part of the teacher, attempts at bribery are perhaps understandable.) Children come to learn the rules of language by being treated *as if* they bear the relevant normative statuses, until eventually they understand enough of the language that they are capable of recognising *on their own* when they do, or do not bear such statuses, and hence are full members of the linguistic community, capable of bearing such statuses in earnest.³⁸

³⁸Plausibly there is no sharp boundary between treating someone *as if* they bear a given normative status and treating them as bearing that status in a full-blooded sense. There is no point at which the relevant statuses simply “descend on the shoulders of the bright child”, as Rorty puts it (Rorty, 1979, p. 187), but rather there is a gradual transition from being an infant who is at best the subject of merely mock attributions to being a full member of the linguistic community who is capable of attributing such statuses to others. This point, and the more general point being made in this paragraph, is indebted to Cavell, 1979, especially pp. 174–176. This paragraph is intended as a corrective of sorts to Sellars’s sometimes overly-behaviouristic description of the process of language learning. (No doubt in part this is a reflection of the influence of Wittgenstein, who unfortunately had a fairly draconian approach to education, both philosophically and practically—cf. Huemer, 2006, p. 207.) Note in connection with this point that I am only claiming that language learning involves *at least* this kind of socially inculcated acquisition of normatively significant pattern-governed behaviour, not that this is *sufficient* for the acquisition of language. Language learning might require inculcating patterns of normatively significant responsive behaviour even if it *also* depends on other things. For example, there are plausibly innate cognitive abilities that both constrain what sorts of languages are learnable *and* play a positive role in making it possible for the language learner to discriminate the (syntactic) features of the linguistic performances to which these normatively significant patterns of behaviour are responses (cf. Chomsky, 1957). Similarly, in the case of perception, there are plausibly innate objectifying perceptual capacities that both constrain the kinds of things that are capable

Coming to be able to see things *as* some way is thus a process of coming to have the ability to recognise certain perceptual situations—namely, those that entitle competent users of the language to claims—*as* situations in which one is entitled to a claim. This process is facilitated by teachers who themselves are able to recognise when the child is in a situation that would, if the child were already a competent language-user herself, confer entitlement upon her. When one is a child, however, and one is still learning, one does not have the ability yet to recognise what it is that one is visually aware of *as* a reason for belief.

There is obviously a sense in which even a child is visually aware of things in her environment, but until she is capable of recognising what conceptual entitlements her experiences confer upon her, she cannot be aware of objects *as* opportunities for judging or applying a concept. Being in a perception that has conceptual content thus depends on having certain conceptual skills: abilities to recognise what it is that one is seeing *as* a reason for applying concepts; abilities that one *may*, but *need not* exercise in order to have an experience with content. It is *this* form of perceptual awareness that confers entitlements upon perceivers, not mere sensory stimulation. As Rorty puts it:

“Sellars invokes the distinction between awareness-as-discriminative-behavior and awareness as what Sellars calls being ‘in the logical space of reasons, of justifying and being able to justify what one says’ [Sellars, 1956/1997, §36]. Awareness in the first sense is manifested by... amoebas and computers; it is simply reliable signaling.³⁹ Awareness in the second sense is manifested only by beings whose behavior we construe as the utterance of sentences with the intention of justifying the utterance of other sentences.” (Rorty, 1979, p. 182)

Rorty’s description here is (unintentionally) apt: Awareness that is normatively significant is *manifested* by beings capable of using language. The perceptual experience itself is a

of being the contents of visual experiences *and* play a positive role in making it possible for the language learner to discriminate the visible features of the environment to which these normatively significant patterns of behaviour are responses, as discussed in chapter 2.

³⁹I have removed Rorty’s claim that rats merely reliably signal, in light of the Burge-inspired criticism I gave of this claim in §2.2.4.

state in which a perceiver is aware of things as potential *opportunities* for applying concepts, and this kind of awareness is *manifested* in the activity of actually applying the concepts that are warranted by one's experience. Being in a perceptual state is being in the logical space of reasons in the sense that in being in a perceptual state one is *able* to justify what one says (or believes) by reference to what one sees, whether or not one actually does this.

Only once a child knows enough of the language to be able to recognise red objects as warranting the application of the concept red will she be capable of actually having the relevant entitlements. The process of coming to be able to be visually aware of a red object in the relevant sense just is the process of coming to be the kind of perceiver who can recognise when she has certain non-inferential entitlements to claims.

One important implication of this story is that unless teachers were able to *recognise* the circumstances in which they ought to attribute perceptual entitlements to a learner, and what entitlements they ought to attribute in those circumstances, language learners could not come to *have* entitlement to claims on the basis of perception. The teacher could not teach the learner to recognise when they have given entitlements unless the teacher herself could recognise when the child has (or will eventually have) these entitlements. But plausibly the *way* that teachers identify the perceptual entitlements that they ought to attribute to language learners is simply by recognising the entitlements that *their* perceptual states confer upon them in the same situation. For learners to be taught to recognise the observational circumstances that confer upon them the relevant normative statuses, and thereby to eventually come to *bear* such normative statuses in earnest, therefore, the features of the learner's experience on the basis of which language users attribute entitlement to such claims must be features that teachers are also capable of recognising as entitling *them* to the very same claims (cf. Davidson, 1990, p. 203; Wittgenstein, 1959/2001, §§ 244, 257).

This is why, above, I described the circumstances of application in terms of the *visible presence of objects to perceivers* rather than, say, in terms of awareness of sense-data or in terms of being in a brain state of a certain kind. If we think of perception as *presenting* a perceiver with mind-independent objects, or entities more generally, then the entities to which the teacher is perceptually related, and on the basis of her relation to which *she* is entitled to claims, are the very same entities to which the *learner* is perceptually related, and on the basis of the *learner's* relation to which the teacher attributes entitlements to her. For example, Benedict, being perceptually related to a white cat, might recognise his entitlement to the claim that the cat is white and, seeing his student also looking at the very same cat, treat him as (or, if he is particularly young, *as if*) he were entitled to this claim *on the very same basis*; namely, on the basis of his *student's* perceptual relation *to the very same cat*.

This story thus constrains the kinds of accounts of perception that are suitable for playing a role in the sort of foundational account of content I gave above. Or, at the very least, it suggests that *if* the story I have told in this section is on the right track then we should think of the features of perceptual states that confer entitlements upon perceivers as things accessible to multiple different perceivers simultaneously. There are, of course, other ways that one might explain how the required kind of convergence between teacher and student is possible without having to endorse a view on which perception directly presents perceivers with mind-independent objects—e.g. if there is sufficient similarity in the brain states of teacher and student that might be enough to get this kind of story off the ground. But this pedagogical story gives us at least a motivation to pursue this kind of account, as I do in the next chapter.

3.6 Conclusion

Wilfrid Sellars famously claimed that:

“in characterizing an episode or a state as that of *knowing*, we are not giving an empirical description of that episode or state; we are placing it in the logical space of reasons, of justifying and being able to justify what one says.” (Sellars, 1956/1997, p. 76)

As Sellars’s justificatory gloss on the idea of the space of reasons indicates, his point is to emphasise that characterising a state as one of knowing is principally to classify it as having a certain normative status; of bearing normative relations to other such states and episodes. The normative pragmatic account of content that I have applied to perception in this chapter generalises this approach. On this view, what it is for a state to be conceptually contentful *at all* is a matter of it having a certain normative role:⁴⁰

“though Sellars here speaks of knowledge in particular, that is just to stress one application of the thought that a normative context is necessary for the idea of being in touch with the world at all, whether knowledgeably or not.” (McDowell, 1996, p. xiv)

Being in an episode or state that counts as *knowing* something requires that the individual have taken a stance on how things are—that they have *undertaken commitment* to a claim. In this chapter I have tried to show that there are benefits both for the philosophy of language and for the philosophy of perception in acknowledging that perceptual states can place perceivers in the space of reasons without those perceivers having to actually judge that things are a certain way.

⁴⁰In light of the discussion in chapter 2, we should be careful to distinguish two interpretations of McDowell’s claim. If we interpret ‘normative context’ in a narrow sense, to mean the kind of linguistically-mediated normativity characteristic of rational creatures then we should restrict this claim to rational creatures. But there is a broader interpretation of ‘normative context’ on which the explanations in which non-rational creatures’ perceptual states feature suffice to supply such a context. In this case, the claim applies to non-rational creatures too. I think the most helpful way to understand the idea, in light of the discussion in the previous chapter, is that like the notion of perception itself, this criterion applies generically to both rational and non-rational animals, but applies in a stricter sense to rational animals.

In the previous chapter I argued that in order for a perceptual state to be attributable to a rational individual as such, that state must be capable of standing in rational relations to that individual's conceptually contentful states, and must be capable of serving as the basis for beliefs. We are now in a position to see the semantic significance of this claim for both perception and language. If a perceptual state is attributable to a perceiver then it *ipso facto* has a normative status sufficient for it to have conceptual content, given the normative pragmatic foundational account of perceptual content I have given above. Thus, given Burge's argument in the previous chapter, for the perceptual states of a rational perceiver to be objective is sufficient for them to have conceptual content. By virtue of the inferential connections that perception has to language, language itself comes to be objectively contentful.

To end this chapter I want to briefly remark on Sellars's suggestion that there is somehow a *contrast* to be drawn between giving an empirical description of a state and placing that state in the logical space of reasons. Rorty describes the purported contrast as follows:

"Sellars's psychological nominalism is not a theory of how the mind works, nor of how knowledge is born in the infant breast, nor of the 'nature of concepts,' nor of any other matter of fact. It is a remark about the difference between facts and rules, a remark to the effect that we can only come under epistemic rules when we have entered the community where the game governed by these rules is played. We may balk at the claim that knowledge, awareness, concepts, language, inference, justification, and the logical space of reasons all descend on the shoulders of the bright child somewhere around the age of four, without having existed in even the most primitive form hitherto. But we do not balk at the thought that a cluster of rights and responsibilities will descend on him on his eighteenth birthday, without having been present in even the most primitive form hitherto... [In] both cases what has happened is a shift in a person's relations with others, not a shift inside the person which now *suits* him to enter into such new relationships." (Rorty, 1979, p. 187)

Rorty is correct that having a certain standing in the space of reasons is not in itself a psychological fact but a normative one. But although this is right, we should not overstate

the independence of the normative and the psychological. Even if, as Rorty suggests, what normative statuses someone has is dependent on their being a member of a (linguistic) community, what *particular* statuses they are capable of bearing in that community will typically *also* depend on their abilities, on what they are capable of *doing*. There is a reason, after all, that we do not let four-year-olds drive cars.

The key to understanding the psychological implications of the claim that perception confers entitlements upon the perceiver is the claim I made in the previous section that a perceiver is only capable of *having* perceptual entitlements once she is capable of *recognising* what entitlements her experience confers upon her.⁴¹ As I put the point in the previous chapter, in order for us to be able to attribute the state *to her*, as a rational creature, that state must not only stand in certain rational relations, but she must be capable of availing herself of some of the opportunities for judgement that her experience confers upon her by *basing beliefs* on her experience.

Fundamentally, therefore, perceiving is being in a certain sensory state in such a way that one is *capable* of recognising what the normative implications of being in that state are. It is, in this sense, a kind of sensory know-how: knowing *how* (or when) to apply concepts to what one sees in a way that will not lead one astray. The things one's senses put one in contact with afford one opportunities for (warranted) judgement, and coming to have contentful experiences is just coming to be able to tell what such opportunities experience affords one.

⁴¹This is not to say that she will always get it right, of course. In order to know a language, and hence for one's utterances to confer linguistic commitments and entitlements on one, a speaker must be capable of recognising what commitments and entitlements oneself and others undertake by asserting sentences of that language, but this does not imply that one cannot be wrong about this in any particular case. The point is just that to be capable of bearing such statuses at all one must be capable of recognising oneself and others as bearing such statuses.

But if a perceiver is capable of recognising what entitlements her experience confers upon her, this must be an aspect of the kind of awareness that perception itself affords her—being in a state that confers entitlements upon one in the way that perception does must also be sufficient for being able to tell what kind of entitlements one bears. After all, how else could she tell, except by means of her experience itself, what kinds of opportunities her experience afforded her? In the next chapter I tackle the issue of how we should think about perception itself such that perception could fulfil both of these roles.

4.0 Perception as a Normatively Significant Relation to Mind-Independent Objects

“[P]erceiving consciousness... *only has to take the object* and to conduct itself as pure apprehension, and what thereby emerges for it is the true. If in this taking, it itself were to do something, it would alter the truth by adding or omitting something... [The perceiver’s] criterion of truth is thus *self-equality*... At the same time, while what is diverse is for the perceiver, the perceiver is a relating of the diverse moments of his comprehending to each other. If in this comparison an inequality differentiates itself, then this is not an untruth of the object, for the object is what is equal to itself; rather, it is an untruth of perceiving.”

—G.W.F. Hegel, *The Phenomenology of Spirit*, §116¹

4.1 Introduction

In the previous chapter, I provided a foundational semantic account of the content of perception that grounded the inferential content of perception in its normative pragmatic role—specifically in the way that perception warrants the applicability of concepts and entitles perceivers to claims. And although the normative pragmatic framework developed in that chapter is *compatible* with a range of views about the metaphysics and phenomenology of perception, I suggested the most *plausible* version of the view pairs this framework with a *relational* view of perception.

On this kind of view, the features of perceptual experience on which the normative features of perception, and hence its content, supervene are *relations to mind-independent objects and their visible properties*.² Combining the normative pragmatic approach to percep-

¹Translation following Pinkard (Hegel, 1807/2018, §116).

²In this chapter I will follow the fairly standard convention of using the term ‘experience’ and the phrase

tual content with this view of the metaphysics and phenomenology of perception would vindicate the (at that point merely provisional) characterisation I gave in the previous chapter of the circumstances of perception that feature in the inferential role of sentences; viz., in terms such as “*the visible presence of a red object to a perceiver*”.

In this chapter I address two issues that face this kind of hybrid account and then draw out the implications of my responses to these problems for our understanding of the content of perception. As we will see there are a growing number of philosophers—usually called *naïve realists*—who endorse views that take relations to environmental entities to be (at least partly) constitutive of the metaphysics and/or phenomenology of perception, as well as giving such relations an important role in the characterisation of the epistemology of perception (§4.2.3). However, most such philosophers advance their views as *rivals* to views that assign content to perception, arguing that the idea that perception is a relation to mind-independent objects is *incompatible* with the claim that perception has content.

The first task of this chapter is thus to argue that there is in fact no conflict between the view that perception has content *as such* and the claim that perception fundamentally consists in a relation to mind-independent objects. Specifically, I argue that although there are some insoluble conflicts between naïve realism and *some* versions of the content view, these conflicts are a consequence not of the claim that perception has content as such but rather of the *foundational* accounts of content that many philosophers give for perception (§4.3.1). By contrast, I argue that not only is there is no conflict between the central metaphysical and phenomenological claims made by naïve realists and the normative pragmatic foundational

‘visual experience’ to refer to both successful (veridical) experiences and hallucinations, and restrict the term ‘perception’ and ‘perceptual experience’ to *successful* experiences (cf. Searle, 1983, p. 38; Crane, 2009, pp. 453–4; Fish, 2009, p. 5). This relational view is thus only a view about *successful* experiences, however, so most of the time I will only talk about such experiences.

account I gave in the previous chapter but that, in fact, if the phenomenological claim made by naïve realists is to play an epistemic role then perception *must* have content (§4.3.2).

The second task of this chapter is to address an objection to this kind of view; namely, I address an argument by Charles Travis that purports to show that even if perception is *associated* with a content, there is no way for the perceiver to be able to *recognise* what the content of her experience is on the basis of purely visual features of her experience alone (§4.4.1). This threatens to undermine the content view altogether, since it is a requirement of this kind of account that this content be conveyed to the perceiver by her experience. I argue that not only can this problem be solved, but that the solution provides a fruitful way for us to think about the relationship between the phenomenological, relational, and conceptual features of experience (§4.4.2).

The third and final task of this chapter is to consider the implications of this kind of hybrid view for our characterisation of the content of perception. I note that if the way things visually look is to make the content of perception recognisable to a perceiver then the way things visually look must itself be conceptually articulable (§4.5.1). To give an account of the content of appearance claims, however, we need to bring in more of the inferentialist semantics than I outlined in the previous chapter (§4.5.2). Once this is in place, I give my account of the content of visual appearances, and explain how this allows us to explain how the inferential content of perception depends on the way that the various things in a perceiver's visual field look, and on the way that these different appearances relate to one another (§4.5.3).

4.2 Naïve Realism and Representationalism

4.2.1 Representationalism in the Philosophy of Perception

In the philosophy of perception, as in the philosophy of language, representational approaches dominate. The orthodox view in contemporary philosophy of perception is *representationalism*; the view that experience represents the world to be a certain way to the perceiver, and accordingly has *content*. To characterise the view in more detail, I will elaborate on the characterisation of representationalism given by Charles Travis:³

1. **Objectivity:** “The representation in question consists in representing things as so (thus, truly/ veridically, or falsely/non-veridically).”
2. **Face Value:** “It has, or gives perceptual experience, a face value, at which it can be taken or declined (or discounted).”
3. **Givenness:** “It is not autorepresentation”⁴
4. **Availability:** “Where we are thus represented to, we can recognize that, and how, this

³The names for these four conditions are not Travis’s, but are given to Travis’s condition by Wilson, 2018, p. 201. I use them for the sake of convenience. It is also worth noting explicitly here that there are two versions of Travis’s paper; the original, published in *Mind* (Travis, 2004), and the reprint, published in the collection *Perception* (Travis, 2013d). Aside from some differences in section numbering and headings, the latter paper also uses different terminology than the former, and includes a few paragraphs of extra discussion not in the original. Given that the terminology in the revised version has some advantages over that used in the original paper, and the extra discussion concerns issues I will address below, I will primarily cite or quote the revised version.

⁴In part of the definition I have excised from this quotation, Travis says that the relevant representing “is allorepresentation, though here, not crucially” (Travis, 2013d, p. 29). I ignore this remark for several reasons. Firstly, Travis’s characterisation of perception as allorepresentation rests, I think, on a highly tendentious application of the distinction between content and its vehicle (Travis, 2013d, p. 26). Although arguing for this claim would take too much space, the tendentiousness of this characterisation is suggested by the fact that unlike the other three claims, which are textually well supported, Travis does not support his attribution of this claim to representationalists with any quotations. Furthermore, Travis hints in this remark, and later explicitly acknowledges, that what really matters is simply that perception is not autorepresentation, not that it is allorepresentation (Travis, 2004, p. 82; Travis, 2018b, p. 341). The essential work that Givenness does for Travis is to allow him to acknowledge that perception might involve or require recognising that various things are so—which is a form of autorepresentation—*without* this constituting an endorsement of representationalism on his part. I discuss this issue more below.

is so; most pertinently, we can appreciate what it is that is thus represented to us as so. Provisionally, I suppose it is (in some sense) the way things look that lets us do that.” (Travis, 2013d, p. 29)

Travis provides a number of examples of philosophers who accept these claims in his paper, but given that they rule out *some* views that could reasonably be characterised as representationalist, I will note potential points at which Travis’s characterisation might diverge from other uses of the term “representationalist” as I discuss each claim.

A good example of the quasi-stipulative nature of these claims concerns Objectivity, which Travis notes is principally included in his characterisation of representationalism to rule out views that claim that perceptual states merely *indicate*—carry information about—features of the environment (Travis, 2013d, p. 24). Although he does not deny that this is a meaningful sense of the term ‘represent’, the representationalists Travis targets have a more substantive conception of representation in mind; representation of something *as so*.

Face Value simply spells out, in slightly idiosyncratic terms, the idea that the content of perception is something a perceiver can *endorse*, though she *need* not do so.⁵ The most straightforward way a representationalist account can satisfy this criterion is for the content of perception to be *propositional*. Propositions are trivially capable of being endorsed because they are capable of being the content of a belief or judgement (Travis, 2013d, p. 25).

However, many representationalists deny that the content of perception is propositional. The majority of representationalists who deny that the content of perception is propositional do so because they claim the content of perception is *nonconceptual* (Crane, 1992b; Dretske, 1981, ch.6; Evans, 1982, p. 227ff; Matthen, 2014; Peacocke, 1989, 1998, 2001).⁶ Peacocke,

⁵On the second part of this claim, see the discussion of Givenness below.

⁶Note that the relevance of this point for Face Value concerns whether perceptual *contents* are nonconceptual, not about whether perceptual *states* are nonconceptual, in Heck’s (2000) sense.

for example, characterises the content of perception in terms of a *scenario*, which is a set of “ways of filling out the space around the perceiver” (Peacocke, 1992, p. 61) with surfaces and their visual properties such as hue, saturation, texture, and so on. Given that scenario contents are not propositions, they are not fit to be the contents of beliefs. However, Peacocke argues that such contents nonetheless determine the conditions under which the perception is accurate or inaccurate, and hence that they may “exclude, or require, the truth of a conceptual content” (Peacocke, 1992, p. 66).⁷ In this sense, therefore, even nonconceptual contents may have a face value.⁸

Not *all* philosophers who deny that perception has propositional content do so because they think the content of perception is nonconceptual, however. On McDowell’s *intuitional* account of content, for example, perception has content akin to that of a demonstrative noun-phrase such as “that red cup” (McDowell, 1998b, 2008a). Although content of this kind is thus not propositional—the phrase “that red cup” does not express a proposition—it nonetheless involves concepts, and so may be endorsed, e.g. by forming the belief that *that red cup is on the table*. Thus, such views also accept Face Value.

Givenness concerns the *way* that contents are represented. Autorepresentation is just *taking something to be so*. The paradigmatic example of autorepresentation is *judgement*—judging that Pangur is a cat *just is* endorsing the claim that Pangur is a cat. Not all autorepresentation need be full-blown endorsement of this kind—taking something to be *doubtfully* or *possibly* so are also forms of autorepresentation according to Travis (Travis, 2004, p. 61fn). Perception cannot be autorepresentation because perception is supposed

⁷Peacocke’s view is more complicated than this, as the location of these features is specified relative to axes that have an orientation determined relative to the perceiver’s body, and only *positioned* scenarios are actually assessable for accuracy (Peacocke, 1992, p. 64), but I pass over the details given that they do not bear on the current point. For the full details, see Peacocke, 1992, ch.3.

⁸Travis does not discuss nonconceptual content explicitly in ‘The Silence of the Senses’, but he does make it clear elsewhere that he takes such views to fall under his purview (Travis, 2009, pp. 165–6; Wilson, 2018, p. 202).

to be a source of information about the world, but to “to autorepresent something is just to accept it” (Travis, 2013d, p. 28). If perception were merely autorepresentation then it would not have a face value that could be endorsed (or not), it would just *be* (a kind of) endorsement.

The meaning of Availability is fairly self evident. The content of perception must be *conveyed to the perceiver* by—or better, *by means of*—her experience, so that the perceiver is capable of accepting or declining this content, in accordance with Face Value (Travis, 2013d, p. 25). This fourth claim thus states that, given that the content must be conveyed to the perceiver by her experience, there must be some way she is able to *tell* what the content of her experience is; some way of *recognising* how things are represented. Some philosophers suggest that there is a weaker claim consistent with representationalism; namely, that perceptual states have contents *associated* with them, even if these are not conveyed to the perceiver herself (Logue, 2014, p. 223). Such a view does not count as a version of representationalism in the relevant sense.⁹

The principal interest of Availability is Travis’s suggestion that the way perceivers recognise the content of perception is on the basis of how things look—a claim he calls ***looks-indexing***. This is important because it plays a crucial role in Travis’s argument against representationalism, which I discuss in §4.4.1. I thus leave the explanation of looks-indexing until later in this chapter. In the next section, I address the question of whether the inferentialist account of perception I have given is a version of representationalism in this sense—that is, whether inferentialists who endorse the account of perception I gave in the previous chapter can, or should, endorse these four claims.

⁹It is worth noting that most representationalists agree. See, e.g., Siegel, 2010, p. 28.

4.2.2 Inferentialism and Representationalism

Given the contrast drawn in the previous chapter between representationalism and inferentialism about the semantics of natural languages, it might be tempting to draw a contrast between inferentialism and representationalism in the philosophy of perception, too. Indeed there are crucial differences, for the responses I give to relationalist criticisms of representationalism below are not available to most representationalists themselves. However, the crucial difference concerns the *foundational* semantic story I told in the previous chapter—the way that contents are grounded directly in the normative significance of perception, rather than, say, (exclusively) in facts about the proper functioning of the visual system, or in facts about the way that the brain processes visual stimuli.¹⁰

Thus, whilst there *are* important differences between inferentialist and representationalist accounts of perception, we need to be careful about how we characterise these differences, so as not to obscure the important similarities. Specifically, in this section I will argue that the inferentialist account of perception satisfies the first three criteria laid down by Travis. (I address the fourth in §4.4.) I begin with the second criterion—the claim that perception has a face value.

On the inferentialist account the content of perception is articulated *fundamentally* in terms of its inferential role rather than in terms of explicitly representational notions like truth-conditions or reference. Despite this difference, however, both inferentialism and representationalism insist, against other accounts such as the sense-data theory and most ver-

¹⁰The parenthetical qualification is necessary given the discussion in chapter 2. That perceptual states are the result of objectifying visual process is *necessary* but not *sufficient* for perception to be attributable to a perceiver. The role that such states play in the constitution of a perceiver's rational perspective is what confers upon these states *conceptual* content, in accordance with the normative pragmatic account of content.

sions of naïve realism, that perceptual states have content.¹¹ As on orthodox representational accounts, this content is fit to serve as the face value of the experience—a perceiver can take her experience at face value by judging or asserting one or more of the claims to which her experience entitles her.

As on orthodox representationalist accounts, however, although perception has content, undergoing an experience with a given content does not require the perceiver to *endorse* that content. This is particularly clear if we cash out the difference between perception and judgement in terms of the foundational semantic account given in the previous chapter: Perception is an state that *confers entitlements* on a perceiver, whereas judgement is an event in which a perceiver *undertakes commitment* to a claim.

There are thus two crucial differences between perception and judgement: Firstly, undertaking a commitment is “*doing something* that licenses or entitles others to attribute it” (Brandom, 1994a, p. 196, my emphasis), whereas perception *confers* entitlements upon a perceiver without her having to *do* anything to warrant this change in normative status.¹² Perception is not something a perceiver *does* but something she *undergoes*. Secondly, in perception a perceiver acquires *entitlements* rather than *commitments*. In order to have an experience with a given content, a perceiver need not endorse or accept the claims to which her experience entitles her, not even “doubtfully”. The inferentialist account thus satisfies Givenness.

We can now turn to Objectivity. Travis’s point in laying down this criterion, as we have seen, is primarily to rule out approaches that treat experiences as states of a perceiver that represent things only in the sense that they carry information about the environment.

¹¹When I speak of ‘inferentialism’ *tout court* in this chapter I almost always mean the inferentialist account of perception and perceptual content given in the previous chapter.

¹²Perceivers typically, if not *necessarily*, issue perceptual judgements in response to their experience. The point here is just to emphasise that their entitlements are not *derived* from their undertaking of these commitments, as the entitlements that follow from asserting are.

The inferentialist account clearly is not just a form of representationalism in this weak sense. However, I want to take this opportunity to address the more substantive question of whether it is right to describe the content of perception as representational in the more substantive sense that applies to Travis’s targets, and hence whether we should describe perception as *representing* that things are the way specified by its content.

In the previous chapter I noted that the fundamental difference between representationalist and inferentialist semantics is one of *explanatory priority* (§3.2.1). Representational approaches to semantics explain the inferential features of meaning and content by reference to representational properties such as truth and reference, rather than taking the former for granted as inferentialism does.¹³ But inferentialists do not *deny* the representational dimension of content, they simply insist that representational locutions must be given an explanation in inferential terms:

“The representational semantic tradition embodies an undeniable insight: whatever is propositionally contentful does necessarily have such a representational aspect; nothing that did not would not be recognizable as expressing a proposition. The point of the inferentialist order of explanation is not to *object* to using representational locutions to talk about semantic content. Inferentialism must be understood instead as a strategy for *understanding* what is said by the use of such locutions.” (Brandom, 1994a, p. 496)

Whilst we do not have the space here to go into the details of Brandom’s execution of this project, we saw one example of this approach in action in the explanation of intentionality in §3.5.2. It is important to note this point because whilst commitment to this order of explanation differentiates inferentialist and representationalist *semantics*, it actually justifies the characterisation of the inferential theory of perception as a version of representationalism *about perception*: The inferential content attributed to perceptual experiences on the infer-

¹³As we saw in the previous chapter, inferentialism takes the notion of good inference as a *semantic primitive*, but this notion is explained *metasemantically* in terms of the inferentialist normative pragmatics (§3.4.1).

entialist view is representational content. Thus, inferentialists, too, can say that perception represents things as so. (And given that there is not a convenient explicitly-inferentialist alternative to the phrase ‘the way perception represents things (as being)’, I will often speak this way below.) Inferentialism only requires that the representational dimension of content must ultimately be explained by reference to its inferential features, rather than taking it for granted as representationalists do.¹⁴ The view thus satisfies Objectivity.

Although we have now dealt with the relationship of inferentialism to the first three of Travis’s criteria, and will not address the fourth until later, there are some terminological and taxonomic issues that need to be addressed before we move on. The issues arise from the fact that although inferentialism is discussed in the philosophy of mind, typically as an alternative view of the content of *belief*, it is not seriously considered as an alternative view of the content of *perception*. Representationalists disagree about how to characterise the content of perception, and on what the correct foundational account of content is, but there is an unquestioned assumption underlying these disagreements; namely, that the most *fundamental* features of perceptual content (and of perception *qua* contentful state or event) are its representational features.¹⁵

Given that, as I have noted, inferentialism, too, aims to explain the representational dimension of content, one might not expect this assumption to be a problem. However, one consequence of the fact that this is an *unquestioned* assumption is that it leads many

¹⁴For the details of the rest of the account, see part two of Brandom, 1994a, especially ch.8; for a shorter overview, see Brandom, 1994b, and for a helpful discussion of some of the terminological issues regarding ‘intentional(ism)’ and ‘representational(ism)’, see Brandom, 1994a, pp. 67–72.

¹⁵The extremely thorough Stanford Encyclopedia article on the contents of perception does not mention inferentialism or even conceptual role semantics once (Siegel, 2016). And although representationalists and relationalists alike often cite Harman’s ‘The Intrinsic Quality of Experience’ (1990) as a canonical example of representationalism (e.g. Travis, 2004, p. 58; Byrne, 2001, p. 201), they inevitably ignore the fact that Harman ultimately cashes out the idea of content in terms of his conceptual role semantics (Harman, 1990, pp. 255–7). Travis’s brief discussion of Harman’s functionalism is the exception that proves the rule (Travis, 2013d, pp. 57–58).

representationalists to conflate the claim that perception has content with other, more specific claims. I mention one important example here—namely, that many representationalists claim that contents just *are* truth- or accuracy-conditions (e.g. Glüer, 2014, p. 77; Siegel, 2010, p. 4). In the context of a discussion of the content of perception that includes inferentialism, however, we should be careful to separate the idea that perception has content from the claim that it has accuracy conditions. Here is how Susanna Schellenberg draws the distinction:

“The accuracy conditions of an experience are often equated with the content of the experience. But this cannot be right. Accuracy conditions need to be distinguished both from the content and from the way the world is. The accuracy conditions of an experience specify the way the world would have to be for the content of the experience to be accurate.” (Schellenberg, 2011b, pp. 725–6)

I will follow Schellenberg in distinguishing the claim that perception has content from the claim that it has truth- or accuracy conditions in this way.¹⁶ I will not assume in what follows that perception has truth- or accuracy conditions, because I think it is a more philosophically substantive task for the inferentialist to derive truth- or accuracy conditions from content than it is for representational semanticists. However, the account I give below and the arguments that I give for it are *compatible* with the claim that perception has truth-conditions (though, given that inferential content is conceptual, not accuracy conditions).

To help avoid confusion, given these complications, I will use the phrase ‘*content view*’ to characterise any view that attributes content to perception, whether that content is characterised fundamentally in terms of its inferential or representational features.¹⁷ Represent-

¹⁶Aside from the reasons mentioned here, such a distinction needs to be made in order to accommodate views that attribute *instructive* contents to perception, which determine *appropriateness* conditions rather than accuracy conditions. See, e.g. Springle, 2019, p. 1211.

¹⁷Both representationalists and relationalists have used the phrase ‘content view’ as a generic name for representationalism (e.g. Brewer, 2011, p. 54; Siegel, 2010, p. 28). My use of this phrase coincides with this more common usage as long as we are careful to separate this out from the more specific characterisations some such authors give of content; e.g. Siegel’s characterisation of contents as accuracy conditions.

ationalism and inferentialism about perception may then be seen as specific ways of cashing out this more basic idea, differing from one another by virtue of the dimension of content they take as *basic*. Thus, whilst I will sometimes speak of perception, on the inferentialist view, as representing things, I will only use the term ‘representationalism’ to refer to versions of the content view that characterise the content of perception fundamentally in terms of its accuracy- or truth-conditions, rather than in terms of its inferential role.

The version of the content view I aim to defend in this chapter, however, is a *hybrid* view that combines an inferentialist view of content with a relational view of the metaphysics and phenomenology of perception. In the next section I therefore explain the core commitments of relational approaches, setting the stage for the account that follows.

4.2.3 The Three Central Claims of Naïve Realism

Although representationalism is still the dominant approach to perception, over the last couple of decades it has come under attack by a number of philosophers who reject this view in favour of *relationalism* or *naïve realism*—the view that perception consists of a *relation* between a perceiver and mind-independent objects.¹⁸ Mike Martin characterises the view as follows:

“The naive realist claims that some sensory experiences are relations to mind-independent objects. That is to say, taking experiences to be episodes or events, the naive realist supposes that some such episodes have as constituents mind-independent objects.” (Martin, 2006, p. 354)

The first central feature of naïve realism, on this way of characterising it, is the insistence on the indispensability of the mind-independent entities—Martin mentions relations

¹⁸The terms ‘relationalism’ and ‘naïve realism’ are usually used interchangeably (e.g. Fish, 2009, p. 3fn). Given that it is important for me to separate out the different roles that relations to mind-independent entities play in the naïve realist theory I will use the name ‘naïve realism’ exclusively to characterise those philosophers who accept all three claims that, I argue, characterise this view.

to objects specifically, but many naïve realists also include relations to properties, events, and other types of entities—to a proper characterisation of the *metaphysics* of perception (Brewer, 2011, pp. 92–3; Logue, 2013, p. 107; Logue, 2014, p. 225; Logue, 2012, p. 211; Martin, 2004, p. 39). I will call this the central *metaphysical thesis* of naïve realism.¹⁹

Martin goes on to say, however, that the fundamental *motivation* for this claim is *phenomenological*: The metaphysical claim is plausible because it best captures the way things seem to the perceiver (cf. Campbell, 2002, pp. 117–120; Martin, 2004, p. 42). The phenomenology of perception is *transparent*—in the sense that introspection on experience “reveals only the mind-independent objects, qualities, and relations that one learns about through perception” (Martin, 2002, p. 378)—and the metaphysical thesis seems to explain this feature of phenomenology well.²⁰ For this reason, many philosophers also characterise the view in phenomenological terms (cf. Brewer, 2011, p. 4; Campbell, 2002, p. 116; Campbell & Cassam, 2014, p. 2; Fish, 2009, pp. 6–15; Gupta, 2019, §115; Soteriou, 2010, pp. 224–5; Soteriou, 2016, p. 59):

“[A]ccording to the naive realist, the phenomenal character of the experience—the property of the experience that types the experience by what it is like to have it—is the property of acquainting the subject with such-and-such a presentational character.” (Fish, 2009, p. 15)

The basic idea behind this phenomenological characterisation of the perceptual relation is that the phenomenology of perception—what it is like for the perceiver to undergo that experience, (e.g.) visually speaking—is constituted by the perceiver’s awareness of the

¹⁹Note that the metaphysical thesis is only a claim about *successful* (perceptual) experiences. Obviously hallucinations are not relations to mind-independent objects. For a discussion of the options available to relationalists (or more broadly, disjunctivists) in dealing with hallucinations, see Byrne and Logue, 2008, and for an extended treatment of the issue, see Fish, 2009, ch.4. I will not discuss hallucinations in this chapter, but Fish’s cognitive-doxastic approach to hallucination fits well with an account of content grounded in its normative role, such as the one defended here.

²⁰Though see Gupta, 2019, §§121–129 for a detailed criticism of the idea that introspection supports this strong statement of transparency and hence, by extension, that the phenomenology of perception is best explained by appeal to the metaphysical thesis.

qualitative properties of the very mind-independent entities to which her experience relates her—what Fish calls the *presentational character* of the visual scene. As Brewer puts it: “The ways things *look* are the ways (perceptually presented) *things* look from that point of view in those circumstances” (Brewer, 2011, p. 99).²¹

As this quotation from Brewer indicates, in order to accommodate situations in which there are differences in the phenomenal character of perception that do not seem to be the result of differences in the mind-independent entities to which a perceiver is related alone, many naïve realists treat the perceiver’s “standpoint” (Campbell, 2009, p. 657)—which includes other general aspects of the viewing conditions such as lighting conditions or the relative orientation of entities to the perceiver—as a “third relatum” (Brewer, 2011, p. 96) that also may partly help determine the phenomenology of perception.

The issue of how to think of the role that the third relatum plays in determining the phenomenology of perception is complicated (cf. French, 2018). For this reason, to sidestep these complications let us borrow a phrase from Fish and call the total set of worldly features a perceiver is aware of, which may include features of the third relatum, “the layout of [her] environment” (Fish, 2009, p. 51). We can then state the second thesis as follows: The central *phenomenological thesis* of naïve realism is the claim that the phenomenal character of perception is constituted solely by a perceiver’s awareness of the layout of her environment.

The third and final aspect of naïve realism I want to discuss is neither metaphysical nor phenomenological, but *epistemic*. For reasons that will become clear, I want to distinguish the epistemic *aspirations* that naïve realists have for their view from the specific claims they make about how it is that perception fulfils these aspirations. Naïve realists typically endorse

²¹I will use the terms ‘phenomenology’ and ‘phenomenal character’ interchangeably. This will not affect the point being made here because insofar as the phenomenology of perception matters to the account that I give below, I will talk principally about how things visually look. However, for a detailed discussion of some of the important terminological and philosophical issues, see Fish’s own discussion in Fish, 2009, ch. 1.

both the metaphysical and the phenomenological thesis because they want to make sense of the idea that by being in a perceptual state, the world itself might come to have some rational bearing on what we are to think (cf. Gupta, 2019, §53; Travis, 2013c, p. 10):

“My basic proposal is that seeing *o* involves conscious acquaintance with *o* itself, the concrete worldly source of the truth that *o* is *F*, in a way that may make it evident to the subject that *o* is an instance of ‘*x* is *F*’ as she understands this, and hence evident that *o* is *F*. Seeing that *o* is *F* is thus a way of its being evident that *o* is *F* and is therefore a way of knowing that *o* is *F*.” (Brewer, 2018, p. 1)

Brewer’s suggestion, though epistemic, makes use of both the metaphysical and the phenomenological thesis: The fact that a perceiver is related to mind-independent objects is what makes it the case that the judgement she bases on her experience might count as *knowledge*: Her judgement counts as knowledge because she is related to the very thing that grounds the truth of the claim that *o* is *F*; namely, *o* itself. In this way the world itself has a rational bearing on what it is correct to think. But the phenomenological thesis plays a crucial role, too, for she is aware of *o* in such a way that it is *evident* to her that *o* is *F*.

To see what kind of epistemic role the phenomenological thesis plays, consider an example: Suppose that Benedict is perceptually aware of a coin that happens to be made of gold, but is covered in a layer of nickel. The coin itself is the “concrete worldly source of the truth” that it is made of gold-plated nickel, and so to that extent Benedict is aware of something that is in fact a reason for judging that the coin is made of gold-plated nickel, in the sense that it is a “consideration that counts in favour of” the belief that the coin is made of gold-plated nickel (Scanlon, 1998, p. 17, quoted in Ginsborg, 2006, p. 287).

But if Benedict merely looks at the coin—as opposed to, say, scratching the surface or running chemical experiments on it—his experience does not thereby *give* him a reason to judge that the coin is made of gold-plated nickel, because it is not *visually* evident to

Benedict that there is gold underneath the nickel. Although the *coin* is a reason to judge that it is made of gold-plated nickel, his *experience* is not a reason for him to judge that it is made of gold-plated nickel, in the sense that his perceptual state is not “a psychological state... in the light of which [his] original belief can be recognized... as rational” (Ginsborg, 2006, p. 290).

The coin itself in fact grounds the truth of all sorts of claims, and so to that extent is a reason for judging all sorts of things, but *which* of those claims Benedict’s experience gives *him* reason to believe depends on how the coin *looks*. Crudely speaking, the metaphysical thesis secures the truth-condition for knowledge but the phenomenological thesis secures the justification condition.

And although Brewer’s proposal is just one way of articulating the central aspirations that naïve realists have for their view, and one way that naïve realists put the first two theses to use, other naïve realists make similar aspirational claims. Travis, for example, says that he wants to understand how perception, by making a perceiver aware of her surroundings, makes her visually aware “of... what settles the question whether P, or makes P likely, or is evidence for P, or is, or would (*ceteris paribus*) be reason to think P” (Travis, 2007, p. 118), and Martin claims that “when a subject is perceiving veridically, then the fact perceived is itself ‘made manifest’ to the subject”, and hence the perceiver “ought to make the judgement that matches what is manifest to him” (Martin, 2002, p. 399).

I will call the idea that by making a perceiver visually aware of the world, perception enables the world to bear rationally, in one or all of the above ways, on what she is to think the central *epistemological aspiration* of naïve realists. Naïve realists, however, endorse a specific view of how we ought to epistemically characterise perception. Specifically, they claim that perception is a form of *conscious acquaintance* with mind-independent entities:

“[For] any property ‘*F*’ of a mind-independent physical object, *o*, it is possible to be visually acquainted with *o* and yet not have visually based perceptual knowledge that *o* is *F*. This is part of the [naïve realist] mantra that conscious acquaintance... is more basic than any relation with facts, propositions, or contents concerning the physical objects in question” (Brewer, 2011, p. 141)

Other naïve realists characterise conscious acquaintance by saying that it is “unanalysable” (Brewer, 2018, p. 2), “primitive” (Campbell, 2002, p. 117), “simple” (Campbell, 2002, p. 115), and “irreducible” (Fish, 2009, p. 14). I will call the claim that perception is a relation of acquaintance more basic than any relation to propositions or content the central *epistemological thesis* of naïve realism (cf. Brewer, 2011, p. 140; Gupta, 2019, §118; Soteriou, 2016, ch.4 §3).

Most naïve realists endorse the epistemological thesis, and hence deny that perception is a state that has content, because they think that the claim that perception has content conflicts with the metaphysical and phenomenological theses, and undermines the epistemic aspirations of naïve realism. I discuss this apparent conflict in the following section.

These three theses are the three central claims made by naïve realists. I have tried to indicate in the explanation given above how these three theses, and the epistemic aspiration of naïve realists, relate to one another. For the purposes of the argument I give below, I will follow naïve realists in assuming that both the metaphysical and phenomenological theses are true and aim to show that although the inferentialist account of perception conflicts with the epistemic thesis, by attributing content to perception, it does not conflict with the epistemic aspirations of naïve realists. In fact, I argue that if perception is to play the epistemic role naïve realists want it to play we should instead endorse a hybrid view on which perception *both* relates a perceiver to mind-independent objects *and* has inferentially-articulated conceptual content.

It is worth noting here, before we move on to the next section, however, that although everything I say in this chapter is compatible with the phenomenological thesis, it is also compatible with a rejection of this claim in favour of a slightly weaker thesis. Anil Gupta argues, I think convincingly, that the naïve realist’s commitment to the phenomenological thesis is driven by their commitment to the epistemic thesis, and hence that if this latter thesis is abandoned—as I argue it should be—one can endorse the metaphysical thesis without being committed to the phenomenological thesis, as on Gupta’s own view (Gupta, 2019, §141; cf. Soteriou, 2016, §4.3).²²

This point is worth mentioning because Gupta also provides good reason to think that the phenomenological thesis, at least in its strongest form, is false (see Gupta, 2019, ch. 5).²³ Ultimately I am more sympathetic with an adverbialist view such as that developed by Mazviita Chirimuuta (2015), on which phenomenology depends *both* on features of mind-independent objects *and* on features of the perceiver. This kind of view presupposes the metaphysical thesis, and explains the phenomenology of perception partly in terms of relations to mind-independent entities, but also allows other factors to partly determine the phenomenology of perception. Below I will assume that the phenomenological thesis is true, however, because I aim to show that my view is compatible with even the strongest phenomenological claims made by relationalists. If I am right in what I have suggested here,

²²Gupta does not endorse the metaphysical thesis exactly as I have presented it here, but rather claims that he will “follow the Naive Realists in recognizing the possibility that ordinary things are *sometimes* presented in experience” (Gupta, 2019, §111, my emphasis). This qualification is important because Gupta has a more liberal conception of presence than naïve realists (cf. Gupta, 2019, §§125–128).

²³Some relationalists seem to accept a weaker claim than this in terms of *determination* rather than *constitution*. The interpretative and philosophical issues regarding *what* the phenomenological commitments of relationalists are are complex, but for some helpful discussions that lay out a number of interpretative options see Gupta, 2019, ch.5, French, 2018, and Logue, 2012. The adverbialist view I express sympathy with here is arguably compatible with weaker interpretations of the phenomenological commitments of naïve realism.

however, then this view will also be compatible with less committal views of phenomenology such as Chirimuuta's, or even Gupta's.

4.3 Inferentialism for Relationalists

4.3.1 Naïve Realism and the Content View

The relationship between the content view and naïve realism is a point of some contention. As I noted above, naïve realism is typically proposed as an *alternative* to representationalist versions of the content view, and most naïve realists argue that the view is incompatible with the claim that perception has content.²⁴ Most often, naïve realists claim that the metaphysical or phenomenological thesis gives the most *fundamental* characterisation of perception, and on this basis argue that if perception has representational content it cannot *most fundamentally* consist in a visual relation to mind-independent entities. Here is how Brewer cashes out the relevant notion of fundamentality (cf. Logue, 2013, p. 107; Logue, 2014, p. 225; Martin, 2004, p. 39):

“The... most fundamental characterization of the nature of perceptual experience... [is] the most fruitful and comprehensive characterization of perceptual experience for the purposes of our overall theoretical understanding of such experience in the philosophy of mind, metaphysics, and epistemology of perception.” (Brewer, 2011, p. 3)

Naïve realists sometimes assume in their arguments against representationalism that the most fruitful characterisation of perception must be given *either* in terms of its content *or*

²⁴Some relationalists *define* relationalism in such a way that it is incompatible with representationalism; for example, by simply *stipulating* that the relation a perceiver bears to her environment is *non-representational* (Martin, 2002, p. 378; Soteriou, 2010, p. 225). Plausibly, this stipulation is primarily intended to differentiate the notion of a relation relevant to the metaphysical thesis from, for example, merely referential relations. The view I defend below is compatible with the claim that perception is a relation in the much more substantial sense articulated by the metaphysical thesis, so I ignore such stipulative exclusions.

in terms of a relation to mind-independent objects. But if this is true, it is not true simply by virtue of the very idea of fundamentality, at least as Brewer defines it.

For one, as Brewer notes, there are a number of different features of perception we might want to explain, as part of explanatory projects in different areas of philosophy. There is thus no reason in principle why the most fruitful characterisation of perception for the purpose of our theoretical understanding in *one* area of philosophy might make essential appeal to the objects to which perception relates a perceiver and the most fruitful characterisation for *another* purpose might make essential appeal to the content of perception (cf. Logue, 2014, pp. 239–240). In which case, the most fundamental characterisation for the purposes of our *overall* theoretical understanding might characterise perception *both* as relational *and* as having content.

A more interesting possibility, however, is that the most fundamental characterisation we can give of experience even *for a single purpose, with respect to a single theoretical domain*, might be one on which perception *both* relates a perceiver to her environment *and* has content. At the very least, if this sort of view is incoherent, it is due to some *specific* conflict between representational and relational features of perception, rather than due to the very notion of fundamentality itself.

So why should we think that there is a conflict between the idea that perception is a state that represents an object as some way and the idea that perception is a state that *relates* a perceiver to an object?

A central reason there is thought to be a conflict between naïve realism and the content view is due to the *foundational accounts of perceptual content* that many representation-
alists give. Many representationalists endorse the view because they hope to provide a naturalistically-acceptable explanation of perception. The phenomenology of perception is

thought to pose a challenge for naturalistic accounts of the mind, given that it is not clear how to explain phenomenology in naturalistically-acceptable terms (cf. Levine, 1983; Nagel, 1974).

Intuitively, however, judgements, beliefs, and many other representational states, *unlike* perceptual experiences, have representational content but no phenomenology. Thus, it is generally thought that there are good prospects for providing a naturalistically acceptable account of representation—a naturalistic *foundational account of content*—in terms of naturalistically intelligible ideas such as information transmission or proper function. A representationalist account of *perception* thus holds out the hope of providing a naturalistically acceptable account of phenomenology, by reducing phenomenology to, or perhaps even identifying it with, the representational properties of perception (cf. Tye, 1992, 1995).

The specific kinds of accounts that representationalists typically give in the service of this naturalistic goal, however, usually lead them to deny both the metaphysical and the phenomenological theses. For example, by giving an account of the metaphysics of perception that can ground this kind of naturalistic account of phenomenology, many representationalists end up treating experiences as brain states, or at least as states purely internal to the perceiver—states, that is, that a perceiver could be in *even if she were not perceptually related to a mind-independent entity*. A good example of this is Tye’s tracking account of representation:

“[For] each state S of object x , within the relevant set of alternative states of x , we may define what the state represents as follows:

S represents that $P =_{df}$ If optimal conditions obtain, S is tokened in x if and only if P and because P ...

The key idea, then, is that representation is a matter of causal covariation or correlation (*tracking*, as I shall often call it) under optimal conditions.” (Tye, 1995, p. 101)

If, as Tye suggests, for an experiential state to represent an object as having a property is for it to *correlate* with that object having that property then perceptual states cannot have mind-independent entities *as constituents*, as we saw that Martin claims they do. To be in a state that merely *correlates* with certain environmental entities is not to be in a state that *fundamentally* consists in a relation to such entities. As Travis puts it: “a representationalist needs a *bearer* of content which... must be [one] which would be present whether I see a donkey or merely experience a ringer for doing so” (Travis, 2013c, p. 10, my emphasis) Thus, in order to provide a naturalistic explanation of what makes perception a representational state—of the features of perception by virtue of which perception has content—representationalists are forced to deny the metaphysical thesis (cf. Travis, 2013a).

Some representationalists even explicitly argue that there is a conflict between naturalistic approaches to representation and the metaphysical thesis. For example, Tyler Burge has argued that the metaphysical thesis is incompatible with an account of perception based on our best scientific understanding of vision, for he argues that empirical approaches to perception are committed to what he calls the *Proximality Principle*:

“On any given occasion, given the total antecedent psychological state of the individual and system, the total proximal input together with internal input into the system suffices to produce a given type of perceptual state, assuming no malfunction or interference.” (Burge, 2005, p. 22)

The Proximality Principle states explicitly that the kind of state one is in when one perceives is determined *solely* by the kind of sensory stimulation one receives (“total proximal input”), plus the condition of one’s body (including one’s psychological states) at the time that one receives that stimulation. If this is true then, as Burge himself argues, perception is fundamentally a kind of state one could be in regardless of whether or not the distal cause of one’s sensory stimulation was actually the mind-independent entities one’s experience rep-

resents or something else. Perception, according to Burge, thus does not not fundamentally consist in a relation to mind-independent entities.²⁵

On the foundational account of perceptual content given in the previous chapter, by contrast, the content of perception is grounded in *normative* features of content, and there is no reason to think that the aspects of perception relevant for understanding its normative significance are limited to events occurring in the brain, or that the normative significance of perception is constrained by something like the Proximality Principle. In fact, as I argued in §3.5.3, given the need for a certain kind of triangulation between teacher and learner in the process of acquiring a language—a process that is necessary for a perceiver to have the normative statuses that ground the content of her experience—the success of the account may even depend on the normative statuses of a perceiver, and hence the content of her experience, being determined by the mind-independent entities to which she is related. There is thus no tension between the inferentialist foundational account of content and the metaphysical thesis, and indeed this account may even require endorsing this thesis.

Secondly, if the phenomenology of perception is determined by the representational properties of experience then the best a representationalist can say is that when perception accurately represents the environment, features of the environment determine the phenomenology of experience by virtue of the way they determine the content of the experience. And although some representationalists claim that this sort of account is sufficient to explain the

²⁵Burge's point is explicitly about *disjunctivism*, which is, roughly, the view that perceptual states differ in kind from (at least) hallucinations (and sometimes also illusory experiences). The three relationalist theses thus each entail respective versions of disjunctivism—metaphysical, phenomenological, and epistemological disjunctivism (see Byrne & Logue, 2008; Haddock & Macpherson, 2008b, for this distinction). Given that relationalism entails disjunctivism, I will only speak of the relevant relationalist theses, rather than introducing a further set of distinctions that would only serve to complicate the matter. In the discussion below, therefore, I often cite claims about, e.g., the incompatibility of representationalism with disjunctivism in connection with discussions of the purported incompatibility of representationalism with *naïve realism*, as I have done here, but given that naïve realism entails disjunctivism there is no harm in this and so I will not note this explicitly in future instances. For responses to Burge's argument, see Campbell, 2010; Goldhaber, 2019; Travis, 2013b.

transparency of perception (Tye, 2002, 2014), naïve realists argue that this falls far short of the idea, made explicit in the phenomenological thesis, that the phenomenal character of perception is constituted solely by the perceiver’s awareness of the layout of her environment (Campbell, 2002, ch. 6).²⁶

It is worth noting that the clash with the phenomenological thesis affects even naturalistically oriented *hybrid* accounts of content, such as Susanna Schellenberg’s view. Schellenberg’s account is based on the idea that in perception perceivers employ concepts, and that as a result of this perception has content corresponding to the concepts employed.²⁷ This is a hybrid account of sorts because employing concepts in this way in some sense relates perceivers to mind-independent objects. However, in her explanation of phenomenology she notes that it is merely the *employing* of conceptual capacities that determines the phenomenology of perception, not whether or not by employing those concepts a perceiver actually succeeds in referring to a mind-independent entity (cf. Schellenberg, 2011a):

“[The] phenomenology of experience corresponds one-to-one with employing concepts in a sensory mode. The sensory modes in question are modes such as seeing, hearing, touching, smelling, and tasting. Since phenomenology corresponds with *employing* concepts rather than with the concepts themselves, whether or not a concept is empty [that is, whether or not anything is actually picked out by that use of the concept] does not have any effect on phenomenology.” (Schellenberg, 2011b, p. 733)

Thus, by attempting to provide a naturalistic *explanation* of phenomenology in terms of its content (or the capacities or brain states that determine this content), representationalists are led to deny the phenomenological thesis, claiming instead that the phenomenology of

²⁶Tye himself often uses the terminology of constitution (e.g. Tye, 2000, p. 48), but this does not demonstrate that he accepts the phenomenal thesis. For more on this point, see the discussions of Tye and Dretske at Chalmers, 2004, p. 156 and Fish, 2009, pp. 10–14.

²⁷Schellenberg has shifted over time towards an even more naturalistic approach to perception, replacing the notion of employing concepts in a sensory mode with “singling out” particulars, which is a “low-level mental capacity” that does not require classification (Schellenberg, 2018, p. 32), but is rather a “proto-conceptual analogue of referring to a particular” (Schellenberg, 2016, p. 36). On the naturalistic credentials of Schellenberg’s account see especially Schellenberg, 2018.

perception is determined by, or perhaps constituted by, features of perception it could have whether or not a perceiver was actually related to mind-independent entities.

By contrast with such accounts, although most inferentialists ultimately want to explain the norms in which content is grounded in naturalistically-acceptable terms (cf. Brandom, 1994a, ch. 1), this account of content clearly is not intended to provide a reductive naturalistic explanation of phenomenology of the kind that many representationalists hope to be able to give. Insofar as one holds out hope that a naturalistic representational explanation of phenomenology can be given, this might seem to be a drawback of the current account. But in the current context it is actually an *advantage*, as it makes this view particularly well-suited for developing a hybrid account of perception, especially compared to other versions of the content view such as those discussed above.²⁸

Given that the account was not motivated by the desire to provide a reductive naturalistic explanation of phenomenology, there is no motivation on this account to argue that the phenomenology of perception is determined by features of perception that it could have whether or not there was really an object present. The view is thus compatible with the naïve realist's account of phenomenology. As was the case with the metaphysical thesis, there may even be distinct *advantages* to be gained by pairing this view of content with the naïve realist's account of phenomenology, given that the inferentialist account is not fit to give a typical reductive representationalist explanation of phenomenology. In fact, I argue below that the way things visually look to a perceiver plays a crucial role in conveying the content of perception to a perceiver (§4.4.2).

²⁸The idea that often naturalistic motives can hamper attempts to explain the normative or epistemic significance of perception is indebted to Gupta's discussion of the distinction between what he calls the *logical* and the *naturalist* inquiries. This idea shows up often throughout *Conscious Experience*—including in criticisms of some views that are direct inspirations for my own account, such as Sellars's (see especially Gupta, 2019, §§ 42–50)—but for Gupta's overview of the distinction, and some of its implications for debates about perception, see Gupta, 2019, §§ 1–8.

There are, thus, good prospects for a hybrid relational account of perception that combines a *normative pragmatic* foundational approach to perceptual content with the metaphysical and phenomenological theses endorsed by naïve realists. In the next section, I give the first part of an *argument* for the claim that perception is both relational and contentful, based on the idea that if perception is to fulfil the naïve realist’s epistemic thesis then it must be a *normatively significant* relation to mind-independent entities.

4.3.2 The Argument from Normative Significance

In the previous section I argued that the fact that the inferentialist account of perception is based on a normative pragmatic foundational account of content makes it more compatible with naïve realism than many representationalist versions of the content view, which are based on foundational accounts that conflict either with the metaphysical thesis or with the phenomenological thesis. In this section I argue that the normative pragmatic foundational account of content actually provides the resources for an *argument* for the claim that perception has content; an argument that is compatible with the naïve realist’s metaphysical and phenomenological thesis.

Recall that the fundamental assumption underlying the foundational account provided in the previous chapter is a commitment to a *normative pragmatic* approach to content:

“[W]hat attributions of semantic contentfulness are *for* is explaining the *normative* significance of [episodes or] states” (Brandom, 1994a, p.143, partly my emphasis)

We saw in the previous chapter how this idea can be applied to give an account of the content of perception in terms of its inferential role, by analogy with the inferentialist’s metasemantic account. The basic thought animating the argument in this section is that if perception is to fulfil the naïve realists’ epistemic aspirations, it must be a *normatively*

significant relation to mind-independent entities. But if perception is a *normatively significant* relation to mind-independent entities this is sufficient to give perception a normative role sufficient for perception to have content, given the normative pragmatic approach to content. If perception is to play a role in the explanation of how *the world* comes to bear on what we are to think, as it must if the naïve realists' epistemic aspirations are to be fulfilled, then we must think of perception as *both* relational *and* contentful (cf. McDowell, 2013a).

The argument itself—what I will call the ***argument from normative significance***—is fairly straightforward, given that most of the heavy lifting is done by the normative pragmatic framework, which we already discussed in detail in the previous chapter. We can summarise this argument as follows:

- P1:** If a rational perceiver is visually aware of mind-independent entities, her experience warrants the application of certain concepts and entitles her to certain claims.
- P2:** If an experience warrants the application of certain concepts and entitles a perceiver to claims, then that experience has conceptual content.
- C:** If a rational perceiver is visually aware of mind-independent entities, then her experience has conceptual content.

(P2) is perhaps the least obvious premise in the argument. However, we have already done the difficult work of arguing for this premise in the previous chapter: (P2) follows straightforwardly from the foundational inferentialist account of perceptual content.

To briefly recap: Perceptual experiences have content in the same way that assertions have content; namely, by virtue of their normative role. What makes this normative pragmatic fact fit to serve as the basis of an account of the *content* of perception is the fact that the entitlements a perceiver acquires are *inferentially articulated*. For example, if (part of) a perceiver's experience entitles her to the claim that *that apple is red*, this same (part of her)

experience will *ipso facto* entitle her to the claim that *that fruit is red*, that *that apple is coloured*, and so on. Experiences, like assertions, “involve alterations of deontic status *that have other alterations of deontic status as their consequences*” (Brandom, 1994a, p. 188, my emphasis). The entitlements that perception confers upon a perceiver are thus sufficient to determine the *inferential role* of her experience, which on an inferentialist account constitutes that experience’s content. Thus, if perception entitles a perceiver to claims then it has a normative significance that is sufficient to confer upon it inferentially-articulated conceptual content.

The only response available to naïve realists is therefore to reject (P1)—to deny, that is, that perception warrants the application of concepts and entitles perceivers to claims. In the rest of this section I consider the prospects for this line of response.

It might seem just *obvious* that naïve realists deny P1, or that they *ought* to. Travis, for example, says that “perception, as such, simply places our surroundings in view; affords us awareness of them. There is no commitment to their *being* one way or another” (Travis, 2013d, p. 33), and claims that “to see the meat on the rug is just to be suitably sensitive, or responsive to it as it then is... To see that the meat is on the rug I must register something else... I must recognize things being as they are as belonging to a certain range of cases” (Travis, 2007, p. 126). And we saw above that Brewer claims that “for any property ‘*F*’ of a mind-independent physical object, *o*, it is possible to be visually acquainted with *o* and yet not have visually based perceptual knowledge that *o* is *F*” (Brewer, 2011, p. 141). I want to argue, however, that most naïve realists do not in fact deny P1, and that they certainly *shouldn’t* if they want to fulfil their epistemic aspirations for perception.

Firstly, note that Travis’s claim that one can *see* a mind-independent object without seeing *that* it is some way is perfectly compatible with P1, given that *seeing that* involves

believing. For similar reasons, a perceiver can be afforded awareness of objects without *the perceiver* being committed to their being one way or another. The foundational inferentialist account of content I gave does not require that a perceiver actually *avail herself* of the relevant warrants or entitlements to be in a state that so warrants or entitles her. In fact, it actually *requires* that perception can confer entitlements upon a perceiver without her undertaking commitment to these claims. If it did not then it would not be *perception* but only the *responses* perception causes that have inferential significance, as on Brandom's view.

For the same reason P1 is also consistent with Brewer's claim, discussed above, that it is always possible for a perceiver to be aware of an object that has a certain property and yet not have *knowledge* that the object has that property. The kind of normatively-significant awareness that a perceiver has in perception does not depend on a perceiver applying a concept or judging that things are the way perception gives her reason to believe they are, so *a fortiori* it is compatible with a perceiver lacking *knowledge*.

Naïve realists *could* still respond to this argument by denying P1. However, given that the connection that this premise draws between the two is purely *normative*, severing this connection comes at a much greater cost than merely insisting that it is possible to be visually aware of something without *exercising* concepts. The naïve realist would have to deny not just that these conceptual responses are *part of* perception but that they are even *warranted* by perception. That is, they would need to claim that it is possible to be visually aware of a mind-independent entity and yet nonetheless fail to have *entitlement* to a claim, or to be *warranted* in applying a concept.

However, most naïve realists *do not* deny that perception has this kind of normative significance. In fact, the explanation they give of the epistemic role of perception often characterises perception itself in normatively significant terms. As we saw in §4.2.3, the

fact that the *objects* a perceiver is related to are reasons for claims in the sense that they make certain claims true is not sufficient to explain the epistemic significance of perception. Rather, an essential part of the account naïve realists give of the epistemic significance of perception is that it makes it *evident* to the perceiver what concepts she ought to apply, or that it makes certain facts *manifest* to her.

A perceiver is not merely related to things that warrant the application of concepts in the sense that applying those concepts to those objects would not lead to a false belief, but the conscious quality of the experience itself makes the perceiver aware of which concepts it is appropriate to apply or which judgements it is appropriate to make. In order to reap the epistemic benefits that the metaphysical thesis makes possible by relating perceivers to truthmakers, therefore, naïve realists have to acknowledge that perception, by *visually* relating perceivers to such entities, has a certain normative significance—it makes the application of concepts or the undertaking of certain commitments “appropriate” or “correct”.

Though we could consider possible responses to these examples, or consider further reasons naïve realists might give for rejecting P1, I want to suggest, on behalf of the naïve realist, that this is in fact not what they ought to object to in this argument. Rather, the fundamental problem with the argument from normative significance is not P1 at all, but rather with what the conclusion purports to actually establish. Recall that above I followed Travis in characterising representationalism about perception by means of four claims: Objectivity, Face Value, Givenness, and Availability. I argued in §4.2.2 that the inferentialist account of perceptual content meets the first three conditions: I argued that perception has conceptual content (Objectivity), that this content is something the perceiver may endorse in judgement (Face Value), but that perceiving is not itself merely just a kind of endorsement (Givenness). But I postponed the question of whether or not the inferentialist account satisfies Availab-

ility until later, and the argument from normative significance does nothing to address this question.

Specifically, as I noted above, an important part of Availability is the idea that the content of perception is not merely *associated* with an experience in some way but is *conveyed* to the perceiver. A perceiver must be in a position to recognise what it is that perception represents to her, e.g. on the basis of how things look. But I have said nothing so far about how it is that a perceiver comes to be aware of what the content of her experience is. Worse still, there is reason to think that the argument I have given, by virtue of the way that inferential content is grounded in the normative statuses of a perceiver, might make it particularly difficult for an account based on this argument to satisfy Availability.

Consider, by analogy, legal entitlements. Turning 18 confers upon someone the entitlement to purchase alcohol in the UK. Someone asleep in their bed at the stroke of midnight of their 18th birthday thereby acquires the entitlement to purchase alcohol, but are not for that reason alone able to tell that they have acquired new entitlement. Turning 18 is thus sufficient to confer this entitlement on someone, but something *else* has to happen for the perceiver to be in a position to know what their new entitlement is—they have to know about British alcohol laws, they have to be awake, they have to know what date it is, they have to know the date of their birthday, they have to know how old they were before the stroke of midnight, and they have to be paying attention to the time.

If perception has content grounded in the way it alters a perceiver's normative entitlements, as the argument from normative significance claims, then the way perception confers entitlements on a perceiver must not be like this: Perception, by relating perceivers visibly to objects, needs not only to confer entitlements on a perceiver, but to do so in such a way that it *ipso facto* puts her in a position to *recognise* what concepts her experience warrants

her in applying, and the claims to which she is entitled. Unless the features of experience that make it contentful are the very same features that allow her to recognise what this content is, her experience could have content that she was not in a position to recognise. The argument from normative significance thus serves to establish that contents are *associated* with perception, but it does not show that these claims are *conveyed* to the perceiver.

Unfortunately for the prospects of the argument from normative significance, and for the attempt to give a hybrid inferentialist view of perception in general, Charles Travis argues that the problem of spelling out how perception conveys contents to a perceiver is one that cannot be solved by *any* theory of perception. Specifically, he argues that there is no way for an account to satisfy both Givenness and Availability. The visual features of perception, Travis claims, fail to determine a specific content for perception, which undermines Availability. However, to give *non*-visual features of perception a role in determining its content would be to invoke autorepresentation, which would undermine Givenness. I discuss this argument in the following section.

4.4 Contentful Visual Awareness

4.4.1 Travis's Argument

In 'The Silence of the Senses', Charles Travis gives an argument that purports to show that there are no features of experience on the basis of which a perceiver might recognise how things are represented to her—that an account that satisfies the first three of his cri-

teria *cannot* satisfy Availability. Travis’s paper is complex, but Keith Wilson provides the following helpful summary of the central argument:²⁹

- P1.** “If visual experiences were p-representational [perceptually representational] then their content would be recognizable in virtue of how, in experience, things perceptually appear, or look (to the subject). (*Looks-indexing*)
- P2.** Visual looks are incapable of making p-representational content recognizable since they are comparative and so equivocal between multiple contents.
- P3.** Thinkable looks are incapable of making p-representational content recognizable since they are not wholly perceptual.
- P4.** There is no further notion of looks that is both wholly perceptual and capable of making p-representational content recognizable.
- C1.** (*From P2 through P4*) The content of visual experiences cannot be recognizable on the basis of how things look (to the subject).
- C2.** (*From P1 and C1*) Visual experiences are not p-representational.” (Wilson, 2018, p. 206)

Travis notes that one *need* not accept P1 in order to satisfy Availability; there could be ways other than looks-indexing by means of which the content of perception might be made recognisable to a perceiver by her experience. I will accept this premise for the sake of argument, however, as I think that the way things look *does* (or can) make content available to a perceiver. Most representationalists have challenged P4 (Byrne, 2009, pp. 438–444; Schellenberg, 2011b, pp. 720–723; Siegel, 2010, pp. 60–63), but Wilson argues persuasively

²⁹Wilson argues that Travis’s argument has been widely misinterpreted, and that all of the most prominent responses by representationalists thus miss the true force of Travis’s argument. Wilson’s summary helps to avoid some of these potential misunderstandings. I think Wilson is right about this—and indeed Travis himself endorses Wilson’s interpretation and his characterisation of the argument (Travis, 2018b, p. 338)—which is why I follow Wilson’s outline in what follows.

that most of these responses miss the point of Travis’s argument (Wilson, 2018, pp. 212–213).

Specifically, many representationalists seem to have misunderstood Travis’s requirement that the content of perception be ‘looks-indexed’ (which is admittedly a somewhat misleading phrase) to mean simply that there is a sense of ‘looks’ on which content *correlates* with how things look. Establishing that there is a connection of some kind between how things look and content does not constitute an answer to Travis’s question of how a perceiver *recognises* what the content of her experience is, for not all uses of the term ‘look’ describe *visual* features of perception.

Travis himself even argues that there is a sense of ‘looks’—*thinkable* looks—that could actually *be* the content of perception (if, *per impossibile*, according to Travis, perception had content), but he does not thereby undermine his own argument because, as P3 states (for reasons I will explain shortly), such contents are not fit to make content *recognisable* (Travis, 2013d, p. 42).

At best, therefore, rejecting P4 might establish, as the argument from normative significance does, that certain contents are *associated* with perception, but it would not answer the question of how a perceiver is able to *recognise* what the content of her experience is. Thus, I will also not challenge P4. The two central premises of this argument for our purposes, therefore, are P2 and P3.

Travis claims that the way things visually look is wholly a matter of what one is perceptually aware of, and so such looks are at least *fit* to make content recognisable. But visual looks cannot *actually* do this because they “point in no one direction” as far as content is concerned (Travis, 2013d, p. 37). In looking the way they visibly do, objects look like any

number of different things, each of which is a potential candidate for the way the experience represents them to be:³⁰

“A peccary, confronted in the right way, may look exactly like a pig (or it may do to us novices). It also, of course, looks just the way a peccary might look (so confronted). It may also look like a tapir, a clever dummy pig, a wax imitation peccary, and so on.” (Travis, 2013d, p. 38).

In being perceptually related to a peccary, I am thus visually aware of how that peccary visibly looks—of its visible properties. The way it visibly looks is the way a peccary might look, so this is one way that the perceptual state might represent things to be; one thing I might take my experience to represent to me given how things visually look. But the way it visibly looks is also the way a *pig*, or a tapir, or any number of other things would look. Each of these are also candidates for the way perception represents things as being, based on how things visually look to the perceiver.

Clearly perception cannot represent *all* of these things as being the case. An experience that represented an animal as a pig *and* a peccary *and* a tapir, etc. would be incoherent (Travis, 2013d, p. 37). Furthermore, Wilson argues that perception cannot simply represent a potentially open-ended *disjunction* of all the things a peccary might be mistaken for, for such contents would be, as Wilson puts it, “effectively self-verifying” (Wilson, 2018, p. 208), which would contradict Objectivity. Thus, if how things visually look makes the content of perception recognisable, there must be something in or about perception itself that identifies one of these ways things might be, consistent with how things visually look, as *the* way

³⁰I will use the phrase ‘how o visibly looks’ to describe o’s perceptible features, and I will use the phrase ‘how things *visually* look’ to describe the way things visually are in experience for a perceiver, phenomenologically speaking, given her awareness of the perceptible features of the things in her environment. The naïve realist’s phenomenological thesis serves to establish this connection. Cf. Brewer’s claim that the “the ways things *look* are the ways (perceptually presented) *things* look from that point of view in those circumstances” (Brewer, 2011, p. 99).

things are represented as being: “one needs a principled way of ignoring some of the specific ways things look, and attending only to others” (Travis, 2013d, p. 37).

The problem for the representationalist, Travis claims, is that nothing else perceptual *can* play this role. The way things visually look cannot select one of these ways rather than another as the way things are represented as being, since *ex hypothesi* the peccary looks the way any of these things *would* look in the right conditions. The only other option, therefore, is that what allows a perceiver to tell how her experience represents things as being is how things *thinkably* look. This brings us into the territory of P3. Here is how Travis describes the notion of thinkable looks:

“The second notion of *look...* is... one on which looking is a matter of some proposition enjoying some status or other *in re* being the thing to think. Looking... on this notion, is a matter of things, or something, having a certain rational force regarding some given proposition; a certain bearing on the thing to think.” (Travis, 2013d, p. 41)

Travis claims that the way things thinkably look cannot make the content of perception *recognisable* because awareness that something is the thing to think, or that a claim, say, is reasonable to believe, is not visual awareness (Travis, 2013d, p. 40). The way things thinkably look could, as I mentioned above, just *be* the content of perception, but it is not the right sort of thing to make it recognisable to a perceiver what this content is:

“[A]wareness that something is the thing to think (or, on some qualification, the qualified thing to think) is not *visual* awareness... For there is nothing which that Pia will sink the putt, or that the painting is a Vermeer, looks like. *That such-and-such* is not the sort of thing to have a look. Unlike, say, the lemon on the counter, it is not the sort of thing to form images—a point of Frege’s.” (Travis, 2013d, p. 40)

Thus, according to Travis, the only kind of awareness a perceiver has of the way things thinkably look is in conceptual *responses* to perception, not in perception itself. Being aware of how the peccary looks to be would simply mean *judging* or otherwise *taking* one of the ways

the peccary looks as the way you think it is. And this, as Travis insists, is just representing things *to oneself* as being that way—for things to thinkably look some way to me is just for them to look *to be* some way to me, and that is “for me so to (auto)represent things to *myself*” (Travis, 2013d, p. 43).

Thus, if the representationalist’s explanation of how a perceiver recognises what the content of her experience is appeals to how things *thinkably* look—how they look *to be* (to the perceiver)—this would mean abandoning Givenness—the claim that perceptual representation is not merely a form of *autorepresentation*, or registration. One’s *deciding* what one thinks the content of perception is cannot be the *basis* on which one recognises what the content of perception is.

Unsurprisingly, given the argument from normative significance, I think that Travis is right to suggest that the way things thinkably look should be thought of as characterising how things are represented to be. Given the normative pragmatic account of content, the content of perception is grounded in the way that perception bears on what to think—that is, in how things thinkably look. But I also think Travis is right to suggest that precisely for this reason how things thinkably look cannot be that by which the content of perception would be made *recognisable*. Below therefore I argue that how things *visually* look actually *can* make the content of experience recognisable to a perceiver. (Given that I will focus primarily on visual looks I will therefore drop the prefix ‘visual’ from now on, unless there is the possibility of confusion.)

Some of the examples that Travis uses to support P2 involve contents one would have to have quite sophisticated conceptual capacities to be able to recognise, such as a painting being a Vermeer (Travis, 2013d, p. 41) or an animal being a peccary (Travis, 2013d, p. 38).

Travis makes it clear, however, that this point applies to more basic visible features, too: “What goes for peccaries goes for something’s being blue” (Travis, 2013d, p. 38).

Whilst ultimately I think it is possible to argue that both sorts of contents can be conveyed to a perceiver by her experience, I also think we need to explain these two kinds of cases differently, for there is an important sense in which contents involving simple visual properties of objects are more basic. Thus, I will argue in the next section that P2 is not true if we restrict our focus to concepts of basic visible properties such as colour or shape, only returning to the question of how a perceiver might be in a position to recognise richer contents in the conclusion of this chapter.

4.4.2 Visual Looks and Visual Context

To give us a point of reference for the discussion below, consider the explanation Travis gives of why visual looks do not ‘point in one direction’ even as regards quite simple representational contents such as *being blue*:

“[A] shirt may look like a blue shirt. In looking as it does, it will also look the way a white shirt would when illuminated in certain ways, or when in certain conditions. It may also look like countless other things. Some such fact might index the representational content of things so looking. But only on condition that the other such facts do not. The problem now is: *what* decides that it is some one such fact, and no other, which plays that role.” (Travis, 2013d, p. 38)

Keith Wilson also notes that similar arguments apply, *mutatis mutandis*, to shape properties (Wilson, 2018, p. 214): A circular surface, such as the surface of a coin, looks elliptical when viewed from certain angles. Thus, the way the surface of a coin looks when I look at it lying on a table whose surface is at, say, a 45-degree angle relative to my line of sight, is a way that a circular thing seen from this angle would look, but it is *also* a way that an

elliptical thing would look when viewed from head-on. So the same problem arises in the case of visible properties such as shape (cf. Noë, 2004, pp. 82–83).

Both Travis and Wilson are right that, *under certain conditions*, differently coloured or differently shaped objects can look the same in respect of colour or shape. But this does not entail, as they claim, that how things look is not fit to index the content of perception, and hence make content recognisable. In terms appropriate to the inferentialist account of content, it does not entail that one cannot tell, on the basis of how things look, that your experience gives you reason to believe that the shirt is blue, or that the surface of the coin is circular. To see why not, consider the following two cases:

Daylight: Bruce, who has perfect eyesight, is looking directly at a pack of Marlboro reds. He is outside, it is the middle of the day, the sky is clear and the sun is shining brightly.

Dive Bar: Bruce, who has perfect eyesight, is looking directly at a pack of Marlboro reds. However, Bruce is currently hanging out in a seedy dive bar and everything he sees is bathed in red light.

Marlboro reds, famously, have a red, white, and black design on their packaging. For ease of reference, let us call the part of the packaging that looks white in ordinary light the *white (part of the) packaging* and the part of the packaging that looks red in ordinary light the *red (part of the) packaging*. And let us stipulate, for the sake of argument, that in Dive Bar the white packaging looks *the very same* shade of red as the red packaging did in Daylight, due to the lighting conditions in the bar. By stipulation, that is, the red packaging in Daylight and the white packaging in Dive bar look the same in respect of colour.³¹

³¹These two claims are not actually equivalent, given that colour is not reducible to hue, but the stipulation does not undermine the argument because we could contrive a more artificial situation in which this claim were literally true, and because acknowledging the differences in luminosity and so on that make this claim implausible would actually make our argument *easier*, since it would give Bruce more visual features to go on in recognising the content of perception. The stipulation is thus a charitable one, dialectically speaking.

Plausibly, given that Bruce is consciously aware of the whole visible environment in both Daylight and Dive Bar, in Daylight his experience entitles him to the claim that the red packaging *is* red, and he can recognise this. But in Dive Bar, although the white packaging *looks* red, his experience *does not* entitle him to the claim that this part of the packing *is* red—because Bruce is aware of the red light that is covering everything in his vicinity. In Daylight, that is, his experience represents the red packaging *as* red, but in Dive Bar, his experience *does not* represent the red packaging *as* red.³²

If looks-indexing is correct then Bruce ought to be able to recognise, based on how things look, that in Daylight his experience of the red packaging represents that part of the packaging *as* red but that in Dive Bar, despite the fact that the white packaging visually looks red, his experience *does not* represent that part of the packaging *as* red. According to Travis, however, this is not possible, because *ex hypothesi* the red packaging in Daylight and the white packaging in Dive Bar look the same in respect of colour.

Considering in isolation Bruce's experience of the red part of the packaging in Daylight, and his experience of the white part of the packaging in Dive Bar, what Travis says is true: These two experiences are both ones in which something visually looks both, or indifferently, red and/or white-in-red-light. So, considering these parts of his experiences in isolation Bruce cannot recognise whether his experience represents things as red or merely as white-in-red-light, or, even worse, perhaps both. But Bruce is *not* considering these two parts of his experiences in isolation. His experiences of these parts of the packaging are just one

³²On a particularly flat-footed conception of the relationship between the phenomenology and content of perception this might not be true, but it is no part of content views *as such* that the localised colour phenomenology of an experience determines the content of perception. If it were true that just because something visually looked some colour F at one part of the visual field then one's experience represented that part of the visual field *as* F then every experience of shadows, reflections, artificial light, and so on would be at least partly inaccurate. Even accounts that claim that the content of perception supervenes on its phenomenology plausibly only mean that it supervenes on the *total* phenomenology of the visual field as a whole. I explain why this should be so in §4.5.3.

part of his broader visual experience of the whole scene. And in particular, Bruce can see the sun shining on the pack of cigarettes in Daylight, and he can see the red light covering everything in Dive Bar. Furthermore, his awareness of the lighting conditions is itself *visual* awareness; a matter of how things *visually* look.

Thus, given that Bruce is visually aware of the good lighting conditions in Daylight, when he sees the red part of the packaging it *does not* look like something white in red light, for it does not look like it is in red light. He is thus capable of recognising, based on how things visually look, that his experience gives him reason to believe that the red part of the packaging *is* red, not white-in-red-light. Thus, he is in a position to recognise, based on how things look, that this claim is part of the inferential content of his experience.

And in Dive Bar, by contrast, given that Bruce is visually aware of the red light covering everything in sight, when Bruce sees the white part of the packaging, it *does not* look like something red in ordinary light, for it does not look like it is in ordinary light. He is thus capable of recognising, based on how things visually look, that his experience *does not* give him reason to believe that the white packaging is red. Thus, he is in a position to recognise, based on how things look, that this claim is *not* part of the inferential content of his experience.

The same kinds of considerations apply, *mutatis mutandis*, to experiences involving shape. When I see a coin lying on a table from the right kind of angle, the surface looks elliptical, and hence in *that* respect visibly looks the way an elliptical object would look if viewed from head-on. But one side *also* looks further away from me than the other; it visually looks tilted away from me. So I am capable of recognising, purely based on how things look, that my experience represents the coin's surface *as* circular, rather than as elliptical.

Thus, the answer to Travis’s question of “*what* decides that is is some one such fact, and no other” that plays the role of allowing the perceiver to recognise what the content of perception is, at least for basic visible properties such as shape and colour: *Other ways things look*. The facts that index the content of perception, allowing the perceiver to recognise what the content of her experience is, are just the various ways other things in the perceiver’s experience visually look; in these cases, her visual awareness of the viewing conditions.

If I see a red apple on a table in good light, therefore, this does not merely entitle me to the claim that there is an apple on the table in a way that has no bearing on the question of whether I can recognise this entitlement—as I might, unbeknownst to me, acquire the entitlement to purchase alcohol upon turning 18 without that event also putting me in a position to recognise that I have this entitlement. Rather, because perception provides *visual* awareness of mind-independent entities it not only confers entitlements on a perceiver, but does so in such a way that she is *ipso facto* able to recognise what these entitlements are, based on how things visually look.

Thus, by undergoing an experience in which things visually look a certain way a perceiver is *ipso facto* in a position to recognise how they look to be and, hence, given the inferentialist’s normative pragmatic account of perceptual content, is in a position to be able to tell what the content of her experience is.

As I noted in §4.2.3, the central reason that naïve realists claim that perceptual relations are more basic than relations to propositions or content is because they think that it is the only way to fulfil their epistemic aspirations for perception. But we are now in a position to see that the hybrid account that I have given also allows the world itself to bear on what to think in just the way I argued that naïve realists claim they want it to: On this hybrid account, perception makes perceivers visually aware of the very things that make their beliefs

reasonable, thus allowing *the world itself*, and not mere intermediaries such as brain states, to bear on what to think.

Furthermore, the *way* that the world bears on what to think is by making a perceiver *visually* aware of things in such a way that the perceiver is in a position to tell, based on how the very layout of her environment visually looks, what her experience gives her reason to believe. Both the metaphysical and the phenomenological thesis thus play a role in this explanation of the epistemic significance of perception, in part because this epistemic explanation is an account of the very same features of perception in which its content is grounded, on the normative pragmatic approach.

This account thus satisfies the epistemic *aspiration* naïve realists have for perception but does so precisely by rejecting the epistemic *thesis*, for given the normative pragmatic account of content I gave in the previous chapter, this means that by undergoing an experience, perception not only has content but visually conveys this content to the perceiver. The fact that perception visually relates perceivers to mind-independent objects is at once both that by virtue of which perception has content—given that this relation is *normatively significant*—and that by which a perceiver is in a position to recognise what this content is—given that by undergoing this experience a perceiver is aware of how things visually look.

Before I move on to the next section, I want to address one potential concern naïve realists might have about the explanation I have just given, a concern that relates to the issue of the epistemic status of perception. This explanation makes essential appeal to the fact that a perceiver has the right kind of conceptual capacities to enable her to judge that things are as her experience gives her reason to believe they are. Her awareness of how things visually look only puts her in a position to recognise how her experience represents things as being because she has certain concepts. In *this* sense, therefore, the way things visually

look *alone* does not allow her to recognise what the content of her experience is—it requires that the perceiver have certain concepts, too.

A naïve realist might object that this violates the requirement that the content of perception be conveyed to a perceiver *purely* by how things visually look. However, I don't think that it is reasonable for Travis, or any other naïve realist, to suggest that what is required to respond to this argument is that perception can “make p-representational content recognizable” (Wilson, 2018, p. 206) to someone *regardless of whether they have the concepts that would enable them to recognise this*. To claim that it this is possible is to fall into the Myth of the Given in just the way that McDowell claims that Travis does:

“Having something Given to one [in the Mythical sense] would be being given something for knowledge without needing to have capacities that would be necessary for one to be able to get to know it. And that is incoherent.” (McDowell, 2008a, p. 1)

If perception makes perceivers aware of how things visually look in such a way that it makes it rational for them to apply certain concepts, and puts them in a position to *recognise* that it is rational to apply certain concepts, then this must require that the perceiver *have* those concepts, on pain of falling into the Myth of the Given.

4.5 The Visual and the Conceptual

4.5.1 Appearances and the Content of Perception

I have argued above that the way things visually look warrants the application of concepts and entitles perceivers to claims, and that it does so in a way she is capable of recognising precisely *as* warranting or entitling her in this way. There is a problem for this view, however, that we need to address. Specifically, if awareness of how the way things visually

look bears on what to think is to count as *rational*, rather than merely causal—and hence for the connection between the way things visually look and the way a perceiver has reason to believe things are is to be a *conceptually contentful* relation—then one must be able to *scrutinise* the reasons that perception purportedly gives one. Here is how McDowell puts this point:

“the faculty of spontaneity carries with it a standing obligation to reflect on the credentials of the putatively rational linkages that, at any time, one takes to govern the active business of adjusting one’s world-view in response to experience.” (McDowell, 1996, p. 40)

As applied to the picture I sketched in the previous section we can put the point as follows: If the application of concepts in response to how things visually look is to be counted as a *rational* response—as it must if the connection between the way things visually look and the way they are taken to be is to be understood as genuinely *conceptual* rather than merely *causal*—then one must be able to question whether the way things visually look *really is* a reason for applying a given concept or not. And McDowell’s point is that in order to be able to do this one must be able to put into words not just one’s conceptual response but also the visual *basis* for this response (cf. Wallage, quoted in Travis, 2018a, p. 50).

In effect, if a perceiver’s conceptual responses to perception stand in relation to the way things visually look as the conclusion of an inference stands to its premise(s), as they do according to the normative pragmatic account of content I have given, then the way things visually look, too, must have conceptual content, and not just the perceiver’s responses to the way things look.³³ To put the point succinctly: The way things visually look can only be a reason to believe something if the way things visually look has conceptual content.

³³This is, of course, not to say that responding to perception by undertaking commitment to a claim is *inferring*. To infer Q from P is to undertake commitment to Q on the basis of one’s commitment to P, but to undertake commitment to a claim in response to perception is to undertake commitment to that claim on the basis of one’s *entitlement* to it.

So far I have been assuming what Kathrin Glüer calls a *naïve semantics* for perception (Glüer, 2009, p. 300). On a naïve semantics for perception, perceptual experiences have as their contents claims of the form *o is F*—claims, that is, that attribute ordinary visible properties like shape and colour to mind-independent objects. If how things visually look could be articulated by a naïve semantics then we could easily satisfy McDowell’s criterion: *That the apple is red* is clearly something that can serve as a reason for one’s belief that the apple is red, and whose rational support for various claims can be scrutinised—e.g. by questioning whether the fact that the apple is red is reason for thinking that it is ripe.

Unfortunately, however, the role I gave to the way things visually look in §4.4.2 precludes such a simple solution to this problem. *If* the way things visually look is sufficiently rich, as in Daylight, then one’s experience will entitle one to a claim of the form *o is F*. However, the argument in §4.4.2 depended on the assumption that the way things visually look does not *always* entitle a perceiver to one to a claim of this form. Specifically, I argued that the way things visually look in Dive Bar does not give him reason to believe that the red packaging is red. Given the normative pragmatic semantics for perception, this means that the this claim is not part of the content of his experience, and so *a fortiori* cannot be the content of how things visually look.

This suggests that we cannot capture the way things visually look to Bruce in Dive Bar in terms of a naïve semantics for perception: Claims of the form *o is F* sometimes *follow* from the way things visually look, but this is a *consequence* of the way things visually look, not a direct *articulation* of the way things visually look. But this brings us back to our original problem: Unless we can articulate the way things visually look *itself*, we cannot

understand this relation of ‘following’ as genuinely *rational*, and hence cannot think of the relations between judgement and perception as genuinely *conceptual*.³⁴

Fortunately, however, the naïve realist account itself suggests a solution to this problem. Recall Brewer’s slogan that “the ways things *look* are the ways (perceptually presented) *things* look from that point of view in those circumstances” (Brewer, 2011, p. 99). The first occurrence of ‘look’ here refers to how things visually look to the perceiver—the visual phenomenology of her experience. But the *second* occurrence of ‘look’ refers to *visible features* of the entities she is visually related to in her experience.

When a perceiver is visually aware of a red object she is aware not *merely* of its redness, but of its *red appearance*; the object *visually* looks red to her because the object itself *visibly* looks red. If the phenomenological thesis is correct then a red object will only visually look red *to a perceiver* if the object itself visibly looks red. But red objects do not always visibly look red; the way they visibly look depends (amongst other things) on the lighting conditions. These are facts about the visible features of objects, not *merely* facts about how things visually look to a given perceiver.

For this reason I propose that in order to properly capture the rational contribution of how things visually look, we should endorse a ***partial phenomenal semantics*** for perception. Glüer characterises a *phenomenal semantics* for perception as follows: “experiences do not ascribe sensible properties like redness or roundness to ordinary material objects, but ‘phenomenal’ ones: properties like looking red or looking round” (Glüer, 2009, p. 298). I

³⁴Note that this observation has a different point than Travis’s observation that visual looks “point in no one direction”. The current problem is not, or at least not *merely*, that Bruce cannot tell what the content of his experience is, but that if we say that his experience does not entitle him to the claim that the wall is red then, given the normative pragmatic account of content, unless his experience entitles him to some other claim *we cannot say that the way things visually look has conceptual content at all*. If this were so then, for the reasons that McDowell mentions, we could not think of his responses to the way things look as rational *even in cases such as Dive Bar*, in which Bruce *can* tell how things are represented to be based on the way things visually look. Unless the way things visually look can itself be articulated conceptually *as such*, one’s awareness of the way things visually look cannot be considered rational.

say that we should endorse a *partial* phenomenal semantics, however, because Glüer claims that appearance properties are represented in experience *rather than* sensible properties like colour or shape properties, and she suggests that *only* appearance properties are represented in perception.

As we will see below, however, I think that in order to capture the content of how things visually look *in full*, we should acknowledge that features *other* than visible appearances can be represented in perception. And I will argue that one consequence of acknowledging this is that by representing phenomenal properties experiences often *ipso facto* represent the intrinsic properties of objects. There is no *conflict* between the claim that perception represents the way things appear and the claim that perception represents some of the ordinary properties of objects. In fact, I will argue, perception represents ordinary properties *by* representing constellations of appearance properties. Hence I will argue that the content of experience is *partially*, but not *totally*, phenomenal, in Glüer's sense.

There are different ways that one can think about the ontological status of appearances—the properties of *looking red* or *looking round* that are represented by phenomenal contents. Alva Noë, for example, argues that in addition to ordinary visible properties such as size and shape, we should acknowledge *perspectival* properties such as size-in-the-visual-field and perspectival shape (Noë, 2004, pp. 82–3). Susanna Schellenberg, similarly, argues that in addition to ordinary visible properties such as size, shape, and colour we need to acknowledge what she calls *situation-dependent properties*, which are properties that objects have relative to certain viewing conditions (Schellenberg, 2008, p. 60).

Both Noë and Schellenberg insist that although these properties are relational, they are nonetheless *objective*, in the sense that “they do not depend on sensations or feelings” (Noë, 2004, p. 83). They are relational properties because they depend on relations between objects

and *other features of the environment*, such as spatial locations (in the case of perspectival properties) or lighting conditions (in the case of situation-dependent colour properties). But they are nonetheless “understood strictly externally... namely, as determined by the object’s intrinsic properties and the situational features” (Schellenberg, 2008, p. 60). In characterising these properties there is thus no need to make reference to the constitution of the *perceiver*, which makes them compatible with, and indeed well suited to a view that endorses, the phenomenal thesis.³⁵

By contrast, Anil Gupta takes a different approach. Gupta defines appearances in terms of a notion of *subjective identity* (Gupta, 2019, § 144–145).³⁶ Appearances, on Gupta’s view, are also relational features of entities presented in experience, but they involve relations to *perceivers* rather than (merely) to other features of the environment (Gupta, 2019, § 148). Which appearances are manifested to a perceiver thus may depend not merely on the features of the entities of which a perceiver is aware, but may also depend on features of the perceiver such as “the state of the subject’s sense organs, the subject’s beliefs and desires, and the subject’s attention and actions” (Gupta, 2019, § 145). Gupta’s view is thus compatible with the metaphysical thesis but not the phenomenological thesis.

The account that I give in the rest of this chapter does not depend on taking a stance on these ontological or phenomenological issues. *All* that matters about appearances, for the account that follows to succeed, is that objects with the same intrinsic properties can manifest different appearances in different viewing conditions, and that the same appearances can be manifestations of different properties under different viewing conditions: In some conditions white objects visibly look red and in others they visibly look white, and something that

³⁵Craig French suggests that positing objective appearance properties of this kind is a plausible, if not the *most* plausible, way to interpret Brewer’s understanding of thin looks (French, 2018, pp. 156–161).

³⁶Note that Gupta distinguishes subjective identity from subjective *indistinguishability* (Gupta, 2019, § 50).

visibly looks red may *be* red but it may also *be* white, depending on the viewing conditions. It is this *structural* feature of appearances, rather than any particular account of their ontology, that is crucial to the semantic account that follows. Partly for this reason, and also to make the discussion itself easier to read, I will simply call the relevant features, whatever their ontological status, *appearances*, and call claims that ascribe appearances to objects *appearance claims*.³⁷

In the next section I lay out the features of inferential semantics that we need to grasp in order to be able to characterise the semantics of appearance claims and, by extension, of the partly phenomenal contents of experience.

4.5.2 Inferential Role and Inferential Significance

To capture the distinctive inferential content of appearance claims, we need to highlight an important feature of inferentialist semantics that we have not discussed so far; namely, that what follows from a claim depends not just on that claim but also on the other claims that one relies on in drawing out its inferential consequences. I noted, when I first introduced inferential semantics back in §3.2.1, that it is crucial to the plausibility of inferentialism as a theory of the meaning of natural languages that the inferential role of a sentence includes not just logically good inferences but also materially good inferences. The inference from “Pangur is a cat” to “Pangur is an animal” is materially good even without the addition of the collateral premise “All cats are animals”.

However, many other materially good inferences *do* depend essentially on collateral premises, even *qua* materially good. For example, there is a materially good inference from

³⁷As I mentioned in the previous chapter, I think that everything I say here is also compatible with an adverbialist approach to phenomenology. On this kind of account appearances would *neither* be properties of objects as such nor properties of experiences, but properties of *perceptual events*.

“This is a blackberry” and “This is green” to “This is unripe”, even though neither “This is a blackberry” nor “This is green” *on their own* imply “This is unripe”. And this inference captures an important part of the material content of the concept blackberry—namely, it captures the fact that green blackberries are unripe.³⁸




Figure 3: The Inferential Role of a Sentence

When we specify the inferential role of a sentence we thus need to specify not just the sets of sentences that follow from that sentence (Λ_i), and the sets of sentences from which it follows (Γ_i), but also the sets of collateral premises, if any, relative to which the premise(s) in question imply the conclusion(s) in question (Δ_i), as in Figure 3.³⁹

The addition of collateral premises does not always *enable* new inferences, as in the example of the green blackberry above, but also sometimes *defeats* old inferences. Not all material inferences are indefeasibly good, as the inference from “Pangur is a cat” to “Pangur

³⁸It is important that we express this by means of a generic rather than as a universally-quantified claim, given that the inference is defeasible—one could have just painted or dyed the blackberry green, or it could be genetically modified. For more on the semantics of generics and their relation to inference, see Stovall, 2019.

³⁹Fundamentally therefore the (nonmonotonic) implication relation holds between *sets* of sentences, rather than individual sentences. Given this fact we could eliminate the sets Δ_i from the left side of this diagram by making the sets of premises from which P follows just the relevant sets $\Gamma_i \cup \Delta_j$. The reason I choose not to do this will become clear once we apply this model to the content of perception. Note that given that Γ_1 may imply P relative to a set of collateral premises Δ_1 but not relative to another set of collateral premises Δ_2 , each set Γ_i and Δ_j may appear more than once on the left hand side of the diagram, as both Γ_1 and Δ_1 do in Figure 3.

is an animal” is. Many, if not *most*, material inferences are only *defeasibly* good (Brandom, 2000, p. 88; Brandom, 2008, pp. 106–107; Sellars, 1953, p. 13). For example, the inference from “Pangur is a cat” to “Pangur has fur” is a good one, in that the premise *entitles* one to the conclusion—that *Pangur is a cat* is a *reason* to believe that Pangur has fur—but one can turn this good inference into a bad one simply by adding extra premises. Thus, the inference from “Pangur is a cat” *and* “Pangur is a Sphynx” to “Pangur has fur” is *not* good.

The role given to collateral premises in the characterisation of inferential roles introduces a modal flavour to the inferentialist account of sentence meaning. A sentence’s inferential role does not actually determine what follows from the sentence, but merely what *would* follow from it *were* one to rely on certain collateral premises in drawing out its consequences. Crudely speaking, you shouldn’t expect the meaning of a sentence to tell you what other premises you ought to use in determining what follows from it. Rather, the role a sentence plays as a premise in inferences is, in effect, a *function* from sets of sentences (collateral premises) to sets of sentences (inferential consequences).⁴⁰ The single-premise inferences in which a sentence features are thus just one *part* of its total inferential role.

Thus, although I have been using the phrases “inferential content” and “inferential role” interchangeably so far, we should, in light of the discussion above, distinguish the inferential role of a sentence, which is functional in the way described above, from its *inferential content*; what, *relative to a given set of collateral commitments*, actually follows from the sentence, and what other sentences it actually follows from:

“if the content of a claim must at least determine what follows from it (what else it commits one to), then since what a claim commits one to depends on what collateral commitments are available to serve as additional premises... the significance of undertaking any particular

⁴⁰This way of putting the point highlights that inferential roles also have a somewhat dynamic flavour. On the dynamic credentials of inferentialism, see Nickel, 2013.

commitment cannot be determined without appeal to... those collateral commitments.”
 (Brandom, 1994a, p. 92)

Thus, consistent with its inferential *role* being the same the very same sentence may, and typically will, have different *contents* relative to different sets of collateral premises. Specifically, the inferential content of a sentence P relative to a set of collateral premises Δ will be the set of sentences that follow from P in conjunction with Δ , and the sets of sentences that, in conjunction with Δ , imply P . We can thus picture the way that a sentence with *the same* inferential role may nonetheless have *different* inferential contents relative to different sets of collateral premises as in Figure 4.⁴¹




Figure 4: Same Inferential Role, Different Inferential Content

It is crucial to note, for present purposes, that collateral premises not only determine what the consequences of application of a sentence are (the sets Λ_i), but also *the appropriate grounds for undertaking commitment to a claim* (the sets Γ_i) (cf. Brandom, 1994a, p. 139).

⁴¹Brandom usually uses the phrase “inferential significance”, or “the significance of acquiring a commitment” (or entitlement), to mean just the *consequences* of a sentence, relative to a set of collateral premises, and often treats these phrases as equivalent to “the content of a claim” or “the content of a commitment”, as he does in the quotation above. What I am calling the inferential content of a claim, depicted in Figure 4, includes *both* the inferential significance of P relative to Δ_i and the appropriate grounds of P relative to Δ_i . However, below I will mostly use the phrases “inferential significance” and “inferential content” interchangeably, as Brandom does, given that experiences only have inferential consequences, not inferential grounds, so (potential) difference in the use of terminology will not matter in what follows.

To put the point in explicitly normative terms, whether entitlement to one (set of) claim(s) entitles a speaker to another claim depends on what other claims she is committed and/or entitled to. Entitlement to the claim that Pangur is a cat entitles one to the claim that Pangur has fur, *unless* one has commitment to the claim that Pangur is a Sphinx. Entitlement to the claim that something is a blackberry may entitle one to the claim that it is unripe, but *only* if one is entitled to the claim that it is green (or to some other claim that also licenses this inference, such as the claim that it is red).⁴²

The same applies, *mutatis mutandis*, for the role that *perceptual* circumstances of application play as appropriate grounds for undertaking commitment to a claim. That is, what claims an experience entitles a perceiver to will depend on what her other commitments and entitlements are. Given that, as I have argued, *perception itself* confers entitlements upon a perceiver, the entitlements that one aspect of her experience confers upon her may thus depend on what entitlements another aspect of her experience confers upon her. It is this feature of inferentialism that will allow us to give a more satisfactory account of the content of perception, and how this content depends on the way things visually look. I turn to this task in the next section.

4.5.3 The Inferential Content of Perception

The problem we are faced with in accounting for the phenomenal content of perception is that something's visually looking red *sometimes*, but not *always*, warrants the application of

⁴²Note the asymmetry between commitments and entitlements here: Commitment to a claim cannot entitle one to new inferences unless one is entitled to that claim. But commitment to a claim *can* defeat one's entitlement to a claim *even if one is not entitled to this claim*. It enough to undermine the rationality of the inference that one *thinks* one is entitled to that claim. This is why I suggested, when I first introduced the notion of inferentialism as a normative use theory of meaning, that the normative differences may make a genuinely *semantic* difference, rather than *merely* being part of the pragmatics on which this semantics is based. If this is right then our diagrams of the inferential significance of a sentence would have to be more complicated than the one I give in Figure 4, but for the sake of simplicity I bracket this complication here.

the concept red. In particular, I argued in §4.4.2 that whether something’s visually looking red warrants the application of the concept red to that thing depends on the other ways things visually look in that experience, and in particular on the perceiver’s visual awareness of the viewing conditions. Given the resources laid out in the previous section, we are now in a position to accommodate this fact semantically.

The key to a proper understanding of the content of looks claims is to acknowledge that whether “o looks red” implies “o is red” depends on what collateral premises one uses to determine its inferential consequences of this first sentence. In conjunction with the collateral premise “o is in good lighting”, “o looks red” does imply “o is red”. But in conjunction with the collateral premise “o is lit by red light”, “o looks red” *does not* imply “o is red”. And this applies just as much to the other *entitlements* that perception confers upon a perceiver as it does to whatever *commitments* a perceiver has.

How things visually look:

Daylight: $\{o_r \text{ looks red}\} \cup \{o_r \text{ is in good lighting}\} \vdash \{o_r \text{ is red}\}$
Dive Bar: $\{o_w \text{ looks red}\} \cup \{o_w \text{ is lit by red light}\} \not\vdash \{o_w \text{ is red}\}$

Figure 5: The Inferential Content of Daylight and Dive Bar

We can exploit the fact that the inferential content of a claim depends on what other claims one uses as collateral premises to capture the way that the inferential consequences of perception depend not just on how individual entities appear, but also on one’s visual awareness of the viewing conditions. Specifically, we can capture this feature of perception if we think of the inferential content of perception not by analogy with a single-premise inference, as I have been assuming so far, but rather by analogy with a *multi-premise* inference.

We can thus capture the way that the contents of Bruce’s two experiences in Daylight and Dive Bar differ as in Figure 5.

In Daylight, the red part of the packaging (o_r) visibly looks red *and* the red part of the packaging is in good lighting conditions. These two premises capture how things visually look to Bruce: Bruce is visually aware both of the red appearance of that part of the packaging and of the good lighting. Together these two premises imply that that part of the packaging *is* red, which is why this claim is part of the inferential content of his experience and, hence, why his experience represents that part of the packaging *as* red.⁴³

By contrast, in Dive Bar, although the white part of the packaging (o_w) also looks red, it is also lit by red light. These two premises capture how things visually look to Bruce: Bruce is visually aware both of the red appearance of that part of the packaging and of the red light covering it and everything else in his visual field. Together these two premises *do not* imply that that part of the packaging *is* red, which is why this claim is *not* part of the inferential content of his experience and, hence, why his experience *does not* represent that part of the packaging *as* red.

In this way, we can capture the connection between the way things look and the way things are represented to be in an experience as a genuinely *conceptual*—indeed, *inferential*—connection. In Daylight Bruce’s experience represents the red part of the packaging *as* red precisely because that part of the packaging looks red and visually looks like it is in good lighting. The way things visually look is thus not just that by which Bruce is able to *recognise* what the content of his experience is, but also that by virtue of which his experience has the specific inferential content that it does.

⁴³The squiggly turnstile indicates a *nonmonotonic consequence relation*. Thus, ‘ $\Gamma \rightsquigarrow \Lambda$ ’ should be read ‘The set of claims Γ implies the set of claims Λ ’, and ‘ $\Gamma \not\rightsquigarrow \Lambda$ ’ should be read ‘The set of claims Γ does not imply the set of claims Λ ’. Note that this does not entail that the inference from Γ to Λ *is* defeasibly good, only that it *may* be defeasibly good. That is, if Γ entails Λ then Γ also implies Λ .

What the actual inferential content of one's experience is—the way one's experience actually represents things to be—thus depends on the way the different features of one's experience bear on one another in something like the way in which what the conclusion of a multi-premise inference is depends on the way that each of the premises bears on what follows from the others. We cannot determine what an experience of a whole visual scene represents by simply *adding together* the various claims that an isolated experience of each part of the scene would represent any more than we can determine what follows from a set of claims by simply *adding together* the various consequences that each of these claims has on its own. “This is a blackberry” does not entail, on its own, “This is unripe”, and “This is green” does not entail, on its own, “This is unripe”, and yet *taken together* the two premises *do* imply “This is unripe”. The inferential content of an experience is determined by the way that diverse range of features a perceiver is visually aware of relate to one another as a unified whole.

Thus, to return to the discussion that led us to reconsider the nature of the content of perception in the first place, we are now in a position to see how the way things visually look can be conceptually articulated *as such*, and hence is capable of exerting a genuinely *rational* constraint on Bruce's responses—on the way he takes them to be, given how things visually look. Given that one can articulate how things visually look, and scrutinise whether the way things visually look really *does* give one reason to believe what one takes oneself to have reason to believe, one can, in McDowell's terms, “reflect on the the credentials of the putatively rational linkages” involved in experience.

Bruce can thus articulate the content of his experience, and *explain* why his experience in Daylight entitled him to the claim that the red part of the packaging is red, by saying “I can tell that that part of the packaging *is* red because it looks red, and the light out here

is good”. In Dive Bar, by contrast, he can explain why his experience in Dive Bar does *not* entitle him to the claim that the white part of the packaging is red, by saying “I can’t tell whether that part of the packaging is red because although it looks red, the light in here is red, so it *might* be red, but it also might just look red because of the light”.

4.6 Conclusion

I have argued that if perception is to play a role in the kind of epistemic account that naïve realists aspire to tell then it must have normative significance and, given the normative pragmatic foundational account of content, thus will have inferentially-articulated conceptual content. And I have also argued that, at least for basic visible properties such as colour and shape, the way things visibly look can make it recognisable to a perceiver what the content of her experience is, and that indeed the particular unity had by one’s visual awareness is what determines what the content of one’s experience is. So I have established that if the metaphysical and the phenomenological theses are true, perception may nonetheless have content that a perceiver is capable of recognising on the basis of visual features of her experience. Furthermore, I argued that this account fulfils the naïve realists’ epistemic aspirations for perception, by allowing the world itself to bear, *visually*, on what to think.

If we accept the suggestion I made in this chapter, however, that perception involves both visual and conceptual awareness, we need not think of the content of perception as exhausted by concepts of basic visible properties such as shape and colour. I want to end this chapter by briefly indicating why not.

Suppose, for instance, that I see what is in fact an E-Type Jaguar, parked at the side of the road. If I stop to admire it, I will undergo an experience in which things visually look

some way to me. For reasons I have explained in this chapter, this experience will warrant the application of certain concepts, such as concepts of colour and shape, and put me in a position to recognise that such concepts are applicable. If we accept the suggestion I made in the previous section, however, that visual looks themselves are conceptually articulable, and hence that what one's experience represents may depend on what other commitments and entitlements one has, we need not think of the content of perception as *exhausted* by concepts of basic visible properties such as shape and colour.

Thus, suppose that I see the E-Type Jaguar parked at the side of the road, and suppose that I know a thing or two about British cars. There is an obvious sense in which I might be in a position to tell, based on how things visually look, that the car I am seeing is an E-Type Jaguar. The elongated bonnet, the distinctive slant of the headlights, the fins that protrude from the front bumper, that particular British Racing shade of green paint. These are all in an obvious sense *visible* properties of the car that we could articulate, perhaps only with some difficulty, in terms of concepts of basic visible properties such as shape, size, location, colour, and so on.

Without a significant amount of background knowledge about British cars, these visible features would not entitle one to the claim that the car is a Jaguar, or to apply the concept E-Type Jaguar in response to the experience. But if one *has* this background knowledge, one's experience can give one *reason* to believe that the car in question is a Jaguar, or at the very least that, *going on how things look*, one has good reason to believe that it is a Jaguar.

There is an obvious sense in which my ability to recognise these more conceptually sophisticated conceptual contents is mediated by my awareness of the more basic visual features of the experience, *in tandem with* with the knowledge that I have about the way that British cars look. A reasonable way to describe the scenario described above, for

example, is that, by making me aware of the visible properties of the car, my experience entitles me to various claims about its basic visual properties, and that these entitlements can combine, as premises, with my background beliefs to entitle me to new claims, in the same kind of way that my experiential entitlements combine to entitle me to new claims in perception itself.

If I'm just learning, and I'm not that skilled at recognising car models then I will have to register these more basic visual features first, and then on the basis of the *beliefs* I form about these visible properties, *infer* from this, along with my background beliefs about what sorts of combinations of visible features E-Types characteristically have, that it is an E-Type.

But it needn't be like this forever. With enough practice, I can recognise *immediately* when I am looking at an E-Type—my perceptual entitlements can combine as premises *directly* with my background beliefs, rather than requiring me to form beliefs on the basis of these perceptual entitlements first. When this happens, it can just be *visually* evident to me that the parked car is an E-Type. After all, just *look* at it! How could it be anything else?

If we adopt the hybrid version of inferentialism I have outlined in this chapter, therefore, we can have our content and see it too.

5.0 Bibliography

- Adams, Z. & Browning, J., eds. *Giving a damn: Essays in dialogue with John Haugeland*. Cambridge (MA): MIT Press, 2016.
- Adams, Z. & Maher, C. (2016). Constancy mechanisms and the normativity of perception. In Z. Adams & J. Browning (Eds.), *Giving a damn: Essays in dialogue with John Haugeland*. Cambridge (MA), MIT Press.
- Alchourrón, C. E., Gärdenfors, P. & Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50(2), 510–530.
- Almog, J., Perry, J. & Wettstein, H., eds. *Themes from Kaplan*. Oxford: Oxford University Press, 1989.
- Augustine. (1991). *Confessions* (H. Chadwick, Trans.). Oxford, Oxford University Press.
- Block, N. (1986). Advertisement for a semantics for psychology. *Midwest Studies in Philosophy*, 10, 615–678.
- Bogdan, R., ed. *Belief: Form, content, and function*. Oxford: Oxford University Press, 1986.
- Boyle, M. (2005). *Kant and the significance of self-consciousness* (Doctoral dissertation). The University of Pittsburgh.
- Boyle, M. (2016). Additive theories of rationality: A critique. *European Journal of Philosophy*, 24(3), 527–555.
- Brandom, R. B. (1983). Asserting. *Noûs*, 17(4), 637–650.
- Brandom, R. B. (1994a). *Making it explicit: Reasoning, representing, and discursive commitment*. Cambridge (MA), Harvard University Press.

- Brandom, R. B. (1994b). Reasoning and representing. In D. Chalmers (Ed.), *Philosophy of mind: Classical and contemporary readings*. Oxford, Oxford University Press.
- Brandom, R. B. (1995). Knowledge and the social articulation of the space of reasons. *Philosophy and Phenomenological Research*, 55(4), 895–908.
- Brandom, R. B. (2000). *Articulating reasons: An introduction to inferentialism*. Cambridge (MA), Harvard University Press.
- Brandom, R. B. (2002a). Non-inferential knowledge, perceptual experience, and secondary qualities: Placing McDowell's empiricism. In N. Smith (Ed.), *Reading McDowell: On Mind and World*. New York, Routledge.
- Brandom, R. B. (2002b). *Tales of the mighty dead: Historical essays in the metaphysics of intentionality*. Cambridge (MA), Harvard University Press.
- Brandom, R. B. (2008). *Between saying and doing: Towards an analytic pragmatism*. Oxford, Oxford University Press.
- Brandom, R. B. (2010). Inferentialism and some of its challenges. In J. Wanderer & B. Weiss (Eds.), *Reading Brandom: On Making it Explicit*. Oxford, Routledge.
- Brandom, R. B. (2015). *From empiricism to expressivism: Brandom reads Sellars*. Cambridge (MA), Harvard University Press.
- Brandom, R. B. (2019). *A spirit of trust: A reading of Hegel's Phenomenology*. Cambridge (MA), Harvard University Press.
- Brandom, R. B. & Scharp, K. (Eds.). (2007). *In the space of reasons: Selected essays of Wilfrid Sellars*. Cambridge (MA), Harvard University Press.
- Brewer, B. (2011). *Perception and its objects*. Oxford, Oxford University Press.
- Brewer, B. (2018). Perceptual experience and empirical reason. *Analytic Philosophy*, 59(1), 1–18.

- Brogaard, B., ed. *Does perception have content?* Oxford: Oxford University Press, 2014.
- Burge, T. (2003). Perceptual entitlement. *Philosophy and Phenomenological Research*, 63(3), 503–548.
- Burge, T. (2005). Disjunctivism and perceptual psychology. *Philosophical Topics*, 33(1), 1–78.
- Burge, T. (2009). Perceptual objectivity. *Philosophical Review*, 118(3), 285–324.
- Burge, T. (2010). *Origins of objectivity*. Oxford, Oxford University Press.
- Burge, T. (2014). Disjunctivism again. *Philosophical Explorations*, 14(1), 43–80.
- Burgess, A. & Sherman, B. (2014a). Introduction: A plea for the metaphysics of meaning. In A. Burgess & B. Sherman (Eds.), *Metasemantics: New essays on the foundations of meaning*. Oxford, Oxford University Press.
- Burgess, A. & Sherman, B., eds. *Metasemantics: New essays on the foundations of meaning*. Oxford: Oxford University Press, 2014.
- Byrne, A. (2001). Intentionalism defended. *The Philosophical Review*, 110(2), 199–240.
- Byrne, A. (2009). Experience and content. *Philosophical Quarterly*, 59(236), 429–451.
- Byrne, A. & Logue, H. (2008). Either/or. In A. Haddock & F. Macpherson (Eds.), *Disjunctivism: Perception, action, knowledge*. Oxford, Oxford University Press.
- Campbell, J. (2002). *Reference and consciousness*. Oxford, Oxford University Press.
- Campbell, J. (2009). Consciousness and reference. In B. P. McLaughlin, A. Beckermann & S. Walter (Eds.), *The Oxford handbook of philosophy of mind*. Oxford, Oxford University Press.

- Campbell, J. (2010). Demonstrative reference, the relational view of experience, and the proximity principle. In R. Jeshion (Ed.), *New essays on singular thought*. Oxford, Oxford University Press.
- Campbell, J. & Cassam, Q. (2014). *Berkeley's puzzle: What does experience teach us?* Oxford, Oxford University Press.
- Cavell, S. (1979). *The claim of reason: Wittgenstein, skepticism, morality, and tragedy*. Oxford, Oxford University Press.
- Chalmers, D., ed. *Philosophy of mind: Classical and contemporary readings*. Oxford: Oxford University Press, 2002.
- Chalmers, D. (2004). The representational character of experience. In B. Leiter (Ed.), *The future for philosophy*. Oxford, Oxford University Press.
- Chirumuuta, M. (2015). *Outside color: Perceptual science and the puzzle of color in philosophy*. Cambridge (MA), MIT Press.
- Chomsky, N. (1957). Review: Some reflections on language games by Wilfrid Sellars. *The Journal of Symbolic Logic*, 22(4), 402–403.
- Collins, J. & Dobler, T., eds. *The philosophy of Charles Travis: Language, thought, and perception*. Oxford: Oxford University Press, 2018.
- Council of Trent. (1941). *Canons and decrees of the Council of Trent* (H. J. Schroeder O.P., Trans.). Rockford, Tan Books.
- Crane, T., ed. *The contents of experience: Essays on perception*. Cambridge: Cambridge University Press, 1992.
- Crane, T. (1992b). The non-conceptual content of experience. In *Aspects of psychologism* (pp. 175–195). Cambridge (MA), Harvard University Press.

- Crane, T. (2006). Is there a perceptual relation? In *Aspects of psychologism* (pp. 196–216). Cambridge (MA), Harvard University Press.
- Crane, T. (2009). Is perception a propositional attitude? *The Philosophical Quarterly*, 59(236), 452–469.
- Crane, T. (2014). *Aspects of psychologism*. Cambridge (MA), Harvard University Press.
- Davidson, D. (1965). Theories of meaning and learnable languages. In *Inquiries into truth and interpretation* (2nd ed., pp. 3–17). Oxford, Oxford University Press.
- Davidson, D. (1967). Truth and meaning. In *Inquiries into truth and interpretation* (2nd ed., pp. 17–37). Oxford, Oxford University Press.
- Davidson, D. (1969). True to the facts. In *Inquiries into truth and interpretation* (2nd ed., pp. 37–55). Oxford, Oxford University Press.
- Davidson, D. (1970). Semantics for natural languages. In *Inquiries into truth and interpretation* (2nd ed., pp. 55–64). Oxford, Oxford University Press.
- Davidson, D. (1973a). In defence of Convention T. In *Inquiries into truth and interpretation* (2nd ed., pp. 65–75). Oxford, Oxford University Press.
- Davidson, D. (1973b). Radical interpretation. In *Inquiries into truth and interpretation* (2nd ed., pp. 125–140). Oxford, Oxford University Press.
- Davidson, D. (1975). Thought and talk. In *Inquiries into truth and interpretation* (2nd ed., pp. 155–171). Oxford, Oxford University Press.
- Davidson, D. (1982). Rational animals. In *Subjective, intersubjective, objective: Philosophical essays, volume 3* (pp. 95–106). Oxford, Clarendon Press.
- Davidson, D. (1988). Epistemology and truth. In *Subjective, intersubjective, objective: Philosophical essays, volume 3* (pp. 177–192). Oxford, Clarendon Press.

- Davidson, D. (1990). Epistemology externalized. In *Subjective, intersubjective, objective: Philosophical essays, volume 3* (pp. 193–204). Oxford, Clarendon Press.
- Davidson, D. (2001a). *Inquiries into truth and interpretation* (2nd ed.). Oxford, Oxford University Press.
- Davidson, D. (2001b). *Subjective, intersubjective, objective: Philosophical essays, volume 3*. Oxford, Clarendon Press.
- Dretske, F. (1981). *Knowledge and the flow of information*. Cambridge (MA), MIT Press.
- Dretske, F. (1986). Misrepresentation. In R. Bogdan (Ed.), *Belief: Form, content, and function*. Oxford, Oxford University Press.
- Dretske, F. (1995). *Naturalizing the mind*. Cambridge (MA), The MIT Press.
- Dummett, M. (1975). What is a theory of meaning? (I). In *The seas of language*. Oxford, Oxford University Press.
- Dummett, M. (1976). What is a theory of meaning? (II). In *The seas of language*. Oxford, Oxford University Press.
- Dummett, M. (1979). What does the appeal to use do for the theory of meaning? In *The seas of language*. Oxford, Oxford University Press.
- Dummett, M. (1993). *The seas of language*. Oxford, Oxford University Press.
- Evans, G. (1982). *The varieties of reference*. Oxford, Oxford University Press.
- Fish, W. (2009). *Perception, hallucination, and illusion*. Oxford, Oxford University Press.
- Fodor, J. (1987). *Psychosemantics*. Cambridge (MA), MIT Press.
- Fodor, J. & Lepore, E. (2001). Brandom's burdens: Compositionality and inferentialism. *Philosophy and Phenomenological Research*, 63(2), 465–481.

- Frege, G. (1879). Begriffsschrift [Excerpts]. In M. Beaney (Ed. & Trans.), *The Frege reader* (M. Beaney, Ed. & Trans.). Oxford, Blackwell.
- Frege, G. (1918). Der Gedanke. In M. Beaney (Ed. & Trans.), *The Frege reader* (M. Beaney, Ed. & Trans.). Oxford, Blackwell.
- Frege, G. (1997). *The Frege reader* (M. Beaney, Ed. & Trans.). Oxford, Blackwell.
- French, C. (2018). Naive realism and diaphaneity. *Proceedings of the Aristotelian Society*, 118(2), 149–175.
- Gamut, L. T. F. (1991). *Logic, language and meaning, Volume 2: Intensional logic and logical grammar*. Chicago, University of Chicago Press.
- Gendler, T. & Hawthorne, J., eds. *Perceptual experience*. Oxford: Oxford University Press, 2006.
- Gentzen, G. (1934a). Untersuchungen über das logische schliessen. I. *Mathematische Zeitschrift*, 36, 176–210.
- Gentzen, G. (1934b). Untersuchungen über das logische schliessen. II. *Mathematische Zeitschrift*, 41, 405–431.
- Gersel, J., Jensen, R. T., Thaning, M. S. & Overgaard, S., eds. *In the light of experience: New essays on perception and reasons*. Oxford: Oxford University Press, 2018.
- Ginsborg, H. (2006). Reasons for belief. *Philosophy and Phenomenological Research*, 72(2), 286–318.
- Ginsborg, H. (2011). Perception, generality, and reasons. In A. Reisner & A. Steglich-Petersen (Eds.), *Reasons for belief*. Cambridge, Cambridge University Press.
- Glüer, K. (2009). In defence of a doxastic account of experience. *Mind & Language*, 24, 297–373.

- Glüer, K. (2014). Looks, reasons, and experiences. In B. Brogaard (Ed.), *Does perception have content?* Oxford, Oxford University Press.
- Goldhaber, C. (2019). Does perceptual psychology rule out disjunctivism in the theory of perception? *Synthese*. <https://doi.org/10.1007/s11229-019-02508-y>
- Gupta, A. (2019). *Conscious experience: A logical inquiry*. Cambridge (MA), Harvard University Press.
- Haddock, A. & Macpherson, F., eds. *Disjunctivism: Perception, action, knowledge*. Oxford: Oxford University Press, 2008.
- Haddock, A. & Macpherson, F. (2008b). Introduction: Varieties of disjunctivism. In A. Haddock & F. Macpherson (Eds.), *Disjunctivism: Perception, action, knowledge*. Oxford, Oxford University Press.
- Harman, G. (1973). *Thought*. Princeton, Princeton University Press.
- Harman, G. (1974). Meaning and semantics. In *Reasoning, meaning, and mind* (pp. 192–205). Oxford, Oxford University Press.
- Harman, G. (1975). Language, thought, and communication. In *Reasoning, meaning, and mind* (pp. 166–182). Oxford, Oxford University Press.
- Harman, G. (1982). Conceptual role semantics. *Notre Dame Journal of Formal Logic*, 23(2), 242–256.
- Harman, G. (1987). (Nonsolipsistic) conceptual role semantics. In *Reasoning, meaning, and mind* (pp. 206–231). Oxford, Oxford University Press.
- Harman, G. (1990). The intrinsic quality of experience. In *Reasoning, meaning, and mind* (pp. 244–261). Oxford, Oxford University Press.
- Harman, G. (1999). *Reasoning, meaning, and mind*. Oxford, Oxford University Press.

- Heck, R. G., Jr. (2000). Nonconceptual content and the “space of reasons”. *The Philosophical Review*, 109(4), 483–523.
- Hegel, G. W. F. (2018). *The phenomenology of spirit* (T. Pinkard, Trans.). Cambridge, Cambridge University Press. (Original work published 1807)
- Heim, I. & Kratzer, A. (1998). *Semantics in generative grammar*. Oxford, Blackwell.
- Horwich, P. (1998). *Meaning*. Oxford, Oxford University Press.
- Horwich, P. (2005). *Reflections on meaning*. Oxford, Oxford University Press.
- Huemer, W. (2006). The transition from causes to norms: Wittgenstein on training. *Gratzer Philosophische Studien*, 71, 205–225.
- Jeshion, R., ed. *New essays on singular thought*. Oxford: Oxford University Press, 2010.
- Joseph, M. (2004). *Donald Davidson*. Chesham, Acumen.
- Kant, I. (2006). *Anthropology from a pragmatic point of view* (R. B. Loudon, Trans.). Cambridge, Cambridge University Press. (Original work published 1798)
- Kaplan, D. (1989). Afterthoughts. In J. Almog, J. Perry & H. Wettstein (Eds.), *Themes from Kaplan*. Oxford, Oxford University Press.
- Khalifa, K., Millson, J. & Risjord, M. (2018). Inference, explanation, and asymmetry. *Synthese*. <https://doi.org/https://doi.org/10.1007/s11229-018-1791-y>
- Kripke, S. (1980). *Naming and necessity*. Cambridge (MA), Harvard University Press.
- Leiter, B., ed. *The future for philosophy*. Oxford: Oxford University Press, 2004.
- Levine, J. (1983). Materialism and qualia: the explanatory gap. *Pacific Philosophical Quarterly*, 64(4), 354–361.
- Lindgaard, J., ed. *John McDowell: Experience, norm, and nature*. Oxford: Blackwell, 2008.

- Logue, H. (2012). Why naive realism? *The Proceedings of the Aristotelian Society*, 112(2), 211–237.
- Logue, H. (2013). Good news for the disjunctivist about (one of) the bad cases. *Philosophy and Phenomenological Research*, 86(1), 105–133.
- Logue, H. (2014). Experiential content and naïve realism: A reconciliation. In B. Brogaard (Ed.), *Does perception have content?* Oxford, Oxford University Press.
- Martin, M. G. F. (2002). The transparency of experience. *Mind and Language*, 17(4), 376–425.
- Martin, M. G. F. (2004). The limits of self awareness. *Philosophical Studies*, 120(1), 37–89.
- Martin, M. G. F. (2006). On being alienated. In T. Gendler & J. Hawthorne (Eds.), *Perceptual experience*. Oxford, Oxford University Press.
- Matthen, M. (1988). Biological functions and perceptual content. *The Journal of Philosophy*, 85(1), 5–27.
- Matthen, M. (2014). Image content. In B. Brogaard (Ed.), *Does perception have content?* Oxford, Oxford University Press.
- McDowell, J. (1995). Knowledge and the internal. *Philosophy and Phenomenological Research*, 55(4), 877–893.
- McDowell, J. (1996). *Mind and world* (2nd ed.). Cambridge (MA), Harvard University Press.
- McDowell, J. (1998a). Intentionality as a relation. In *Having the world in view: Essays on Kant, Hegel, and Sellars*. Cambridge (MA), Harvard University Press.
- McDowell, J. (1998b). The logical form of an intuition. In *Having the world in view: Essays on Kant, Hegel, and Sellars*. Cambridge (MA), Harvard University Press.

- McDowell, J. (1998c). Sellars on perceptual experience. In *Having the world in view: Essays on Kant, Hegel, and Sellars*. Cambridge (MA), Harvard University Press.
- McDowell, J. (2002). Knowledge and the internal revisited. In *The engaged intellect* (pp. 279–287). Cambridge (MA), Harvard University Press.
- McDowell, J. (2008a). Avoiding the myth of the given. In J. Lindgaard (Ed.), *John McDowell: Experience, norm, and nature*. Oxford, Blackwell.
- McDowell, J. (2008b). Reply to Kathrin Glüer. *Theoria*, 70(2–3), 213–215.
- McDowell, J. (2009a). *The engaged intellect*. Cambridge (MA), Harvard University Press.
- McDowell, J. (2009b). *Having the world in view: Essays on Kant, Hegel, and Sellars*. Cambridge (MA), Harvard University Press.
- McDowell, J. (2011a). *Perception as a capacity for knowledge*. Milwaukee, Marquette University Press.
- McDowell, J. (2011b). Tyler Burge on disjunctivism. *Philosophical Explorations*, 13(3), 243–255.
- McDowell, J. (2013a). Perceptual experience: both relational and contentful. *European Journal of Philosophy*, 21(1), 144–157.
- McDowell, J. (2013b). Tyler Burge on disjunctivism (II). *Philosophical Explorations*, 16(3), 259–279.
- McLaughlin, B. P., Beckermann, A. & Walter, S., eds. *The Oxford handbook of philosophy of mind*. Oxford: Oxford University Press, 2009.
- Millikan, R. (1984). *Language, thought, and other biological categories*. Cambridge (MA), MIT Press.
- Millikan, R. (1989). Biosemantics. *Journal of Philosophy*, 86(6), 281–297.

- Motörhead. (2002). The game [Song]. On *Hammered* [Album]. London, Sanctuary Records.
- Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review*, 83(4), 435–450.
- Nickel, B. (2013). Dynamics, Brandom-style. *Philosophical Studies*, 162(2), 333–354.
- Noë, A. (2004). *Action in perception*. Cambridge (MA), MIT Press.
- Nudds, M. (2012). Origins of objectivity. *Analysis*, 72(1), 157–174.
- Papineau, D. (1984). Representation and explanation. *Philosophy of Science*, 51(4), 550–572.
- Peacocke, C. (1989). Perceptual content. In J. Almog, J. Perry & H. Wettstein (Eds.), *Themes from Kaplan*. Oxford, Oxford University Press.
- Peacocke, C. (1992). *A study of concepts*. Cambridge (MA), The MIT Press.
- Peacocke, C. (1998). Nonconceptual content defended. *Philosophy and Phenomenological Research*, 58(2), 381–388.
- Peacocke, C. (2001). Does perception have nonconceptual content? *The Journal of Philosophy*, 98(5), 239–264.
- Peacocke, C. (2014). Perception, biology, action, and knowledge. *Philosophy and Phenomenological Research*, 88(2), 477–484.
- Peregrin, J. (2006). Meaning as an inferential role. *Erkenntnis*, (64), 1–35.
- Peregrin, J. (2014). *Inferentialism: Why rules matter*. New York, Palgrave Macmillan.
- Radford, L. (2017). On inferentialism. *Mathematics Education Research Journal*, 29, 493–508.
- Reisner, A. & Steglich-Petersen, A., eds. *Reasons for belief*. Cambridge: Cambridge University Press, 2011.
- Rorty, R. (1979). *Philosophy and the mirror of nature*. Princeton, Princeton University Press.

- Russell, B. (1905). On denoting. *Mind*, 14(56), 479–493.
- Russell, B. (1972). *The problems of philosophy*. Oxford, Oxford University Press. (Original work published 1912)
- Ryle, G. (1949). *The concept of mind* (60th Anniversary Edition). New York, Routledge.
- Scanlon, T. M. (1998). *What we owe to each other*. Cambridge (MA), Harvard University Press.
- Schellenberg, S. (2008). The situation-dependency of perception. *The Journal of Philosophy*, 105(2), 55–84.
- Schellenberg, S. (2011a). Ontological minimalism about phenomenology. *Philosophy and Phenomenological Research*, 83(1), 1–40.
- Schellenberg, S. (2011b). Perceptual content defended. *Noûs*, 45(4), 714–750.
- Schellenberg, S. (2016). Perceptual particularity. *Philosophy and Phenomenological Research*, 93(1), 25–54.
- Schellenberg, S. (2018). *The unity of perception: Content, consciousness, evidence*. Oxford, Oxford University Press.
- Searle, J. R. (1983). *Intentionality: An essay in the philosophy of mind*. Cambridge, Cambridge University Press.
- Sellars, W. (1953). Inference and meaning. In R. B. Brandom & K. Scharp (Eds.), *In the space of reasons: Selected essays of Wilfrid Sellars* (pp. 3–27). Cambridge (MA), Harvard University Press.
- Sellars, W. (1954). Some reflections on language games. In R. B. Brandom & K. Scharp (Eds.), *In the space of reasons: Selected essays of Wilfrid Sellars* (pp. 28–56). Cambridge (MA), Harvard University Press.

- Sellars, W. (1963). Phenomenalism. In R. B. Brandom & K. Scharp (Eds.), *In the space of reasons: Selected essays of Wilfrid Sellars*. Cambridge (MA), Harvard University Press.
- Sellars, W. (1968). *Science and metaphysics: Variations on Kantian themes*. New York, Humanities Press.
- Sellars, W. (1969). Language as thought and communication. In R. B. Brandom & K. Scharp (Eds.), *In the space of reasons: Selected essays of Wilfrid Sellars* (pp. 57–80). Cambridge (MA), Harvard University Press.
- Sellars, W. (1970). Toward a theory of the categories. In J. F. Sicha (Ed.), *Kant's transcendental metaphysics: Sellars' Cassirer lectures notes and other essays* (pp. 474–491). Atascadero, Ridgeview.
- Sellars, W. (1974). Meaning as functional classification. In R. B. Brandom & K. Scharp (Eds.), *In the space of reasons: Selected essays of Wilfrid Sellars* (pp. 81–100). Cambridge (MA), Harvard University Press.
- Sellars, W. (1981). The lever of Archimedes. In R. B. Brandom & K. Scharp (Eds.), *In the space of reasons: Selected essays of Wilfrid Sellars* (pp. 229–257). Cambridge (MA), Harvard University Press.
- Sellars, W. (1997). *Empiricism and the philosophy of mind*. Cambridge (MA), Harvard University Press. (Original work published 1956)
- Sellars, W. (2002). *Kant's transcendental metaphysics: Sellars' Cassirer lectures notes and other essays*. Atascadero, Ridgeview.
- Siegel, S. (2010). *The contents of visual experience*. Oxford, Oxford University Press.
- Siegel, S. (2016). The contents of perception [Winter Edition]. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Winter Edition. Metaphysics Research Lab, Stan-

- ford University. <https://plato.stanford.edu/archives/win2016/entries/perception-contents/>
- Simonelli, R. (forthcoming). The normative/agentive correspondence. *The Journal of Transcendental Philosophy*.
- Smith, N., ed. *Reading McDowell: On Mind and World*. New York: Routledge, 2002.
- Soteriou, M. (2010). Perceiving events. *Philosophical Explorations*, 13(3), 223–241.
- Soteriou, M. (2016). *Disjunctivism*. Oxford, Routledge.
- Speaks, J. (2019). Theories of meaning [Winter Edition]. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Winter Edition. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2019/entries/meaning/>
- Springle, A. A. (2019). Perception, representation, realism, and function. *Philosophy of Science*, 86(5), 1202–1213.
- Stalnaker, R. (1997). Reference and necessity. In *Ways a world might be: Metaphysical and anti-metaphysical essays* (pp. 165–188). Oxford, Oxford University Press.
- Stalnaker, R. (2003). *Ways a world might be: Metaphysical and anti-metaphysical essays*. Oxford, Oxford University Press.
- Stovall, P. (2019). Characterizing generics are material inference tickets: A proof-theoretic analysis. *Inquiry*. <https://doi.org/10.1080/0020174X.2019.1580839>
- Strawson, P. F. (1959). *Individuals: An essay in descriptive metaphysics*. London, Methuen.
- Strawson, P. F. (1966). *The bounds of sense: An essay on Kant's Critique of Pure Reason*. London, Methuen.
- Taylor, H. (forthcoming). Fuzziness in the mind: Can perception be unconscious? *Philosophy and Phenomenological Research*. <https://doi.org/10.1111/phpr.12592>

- Travis, C. (2004). The silence of the senses. *Mind*, 113(449), 57–94.
- Travis, C. (2007). Reason’s reach. In *Perception: Essays after Frege* (pp. 118–143). Oxford, Oxford University Press.
- Travis, C. (2009). The inward turn. In *Perception: Essays after Frege* (pp. 144–177). Oxford, Oxford University Press.
- Travis, C. (2013a). Appendix to the preserve of thinkers. In *Perception: Essays after Frege* (pp. 352–363). Oxford, Oxford University Press.
- Travis, C. (2013b). Desperately seeking psi. In *Perception: Essays after Frege* (pp. 259–312). Oxford, Oxford University Press. Revised version.
- Travis, C. (2013c). *Perception: Essays after Frege*. Oxford, Oxford University Press.
- Travis, C. (2013d). The silence of the senses. In *Perception: Essays after Frege*. Oxford, Oxford University Press. Revised version.
- Travis, C. (2018a). The move, the divide, the myth, and its dogma. In J. Gersel, R. T. Jensen, M. S. Thaning & S. Overgaard (Eds.), *In the light of experience: New essays on perception and reasons*. Oxford, Oxford University Press.
- Travis, C. (2018b). Reply to Keith A. Wilson. In J. Collins & T. Dobler (Eds.), *The philosophy of Charles Travis: Language, thought, and perception*. Oxford, Oxford University Press.
- Tye, M. (1992). Visual qualia and visual content. In T. Crane (Ed.), *The contents of experience: Essays on perception*. Cambridge, Cambridge University Press.
- Tye, M. (1995). *Ten problems of consciousness: A representational theory of the phenomenal mind*. Cambridge (MA), The MIT Press.
- Tye, M. (2000). *Consciousness, color, and content*. Cambridge (MA), The MIT Press.

- Tye, M. (2002). Representationalism and the transparency of experience. *Nôus*, 36(1), 136–151.
- Tye, M. (2014). Transparency, qualia realism, and representationalism. *Philosophical Studies*, 170(1), 39–57.
- Vendler, Z. (1957). Verbs and times. *The Philosophical Review*, 66(2), 143–160.
- Wanderer, J. (2008). *Robert Brandom*. Oxford, Routledge.
- Wanderer, J. & Weiss, B., eds. *Reading Brandom: On Making it Explicit*. Oxford: Routledge, 2010.
- Warren, M. D. (2015). Moral inferentialism and the Frege-Geach problem. *Philosophical Studies*, 172(11), 2859–2885.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford, Oxford University Press.
- Wilson, K. (2018). Are the senses silent? Travis’s argument from looks. In J. Collins & T. Dobler (Eds.), *The philosophy of Charles Travis: Language, thought, and perception*. Oxford, Oxford University Press.
- Wittgenstein, L. (2001). *Philosophical investigations* (G. E. M. Anscombe, Trans.; 3rd ed.). (Original work published 1959)