

Andrea Giananti

Openness to the World: an Enquiry into the Intentionality of Perception

Thèse de Doctorat présentée devant la Faculté des Lettres de l'Université de Fribourg, en Suisse.

Approuvé par la Faculté des Lettres sur proposition des professeurs Gianfranco Soldati (premier rapporteur), Bill Brewer (deuxième rapporteur) et Heather Logue (troisième rapporteuse).

Fribourg, le 25 septembre 2015. La Doyenne Prof. Bernadette Charlier Pasquier.

Table of Contents

Acknowledgments (p. 4)

Abstract (p. 6)

Introduction (p. 8)

Chapter One: Intentionality (p. 13)

Abstract (p. 13)

Introduction (p. 13)

1. Objects (p. 15)

2. Content (p. 19)

2.1 Non-Existent Objects (p. 19)

2.2 Modes of Presentation (p. 21)

3. Satisfaction Conditions, Intentionality and Intensionality-with-an-s (p. 25)

4. Perception (p. 29)

4.1 Perception and Intentionality (p. 29)

4.2 Openness to the World (p. 32)

Chapter Two: The Way Things Look: a Defence of Content (p. 39)

Abstract (p. 39)

Introduction (p. 39)

1. The Silence of the Senses (p. 41)

1.1 Setting the Stage (p. 41)

1.2 A Dilemma for the Representationalist (p. 43)

1.3 Phenomenal Looks (p. 49)

1.4 Content and Accessibility (p. 56)

2. The Way Things Look (p. 68)

2.1 Perceptual Content and Perceptual Constancies (p. 68)

2.2 Perceptual Content, Adverbial Modifications and Dynamic Perception (p. 71)

2.3 A Sense of the World: Content and Mind-Independent Objects (p. 75)

2.4 Objections and Replies (p. 79)

Chapter Three: Perception and Illusion (p. 84)

Abstract (p. 84)

Introduction (p. 84)

1. Sense-Data (p. 86)

2. Representationalism (p. 90)

3. Pure Relationalism, Naïve Realism and Disjunctivism (p. 104)

4. Pure Relationalism and Perceptual Error (p. 113)

4.1 Mixed Accounts (p. 118)

4.2 Appearance Properties (p. 120)

4.3.1 Visual Similarities (p. 122)

4.3.2 Phenomenological Intelligibility and Sensible Qualities (p. 127)

Chapter Four: Perception, Phenomenal Character and Particularity (p. 138)

Abstract (p. 138)

Introduction (p. 138)

1. A Representationalist Argument (p. 143)

2. Particularity and Generality (p. 148)

3. Perception and Generality (p. 153)

3.1 Perception, Determinacy and Similarities (p. 155)

3.2 A Two-Headed Monster and the Generality of Perception (p. 158)

Conclusion: Reconciling Naïve Realism and Content? (p. 164)

References (p. 172)

Acknowledgements

Many people have contributed in several ways to making this thesis better than it would have been otherwise.

I could not have started this work without my main supervisor, Gianfranco Soldati. Gianfranco encouraged me, back in 2009, to write the research proposal that later became the starting point for this thesis, and he has been stimulating and supportive ever since. My numerous and long philosophical discussions with him taught me always to see the contemporary literature within a larger philosophical context, and the present work is informed through and through by ideas that emerged during those discussions.

Bill Brewer supervised me during my time as a visiting doctoral student at King's College London. My chats with him helped me greatly to better understand fundamental questions in the philosophy of perception, and his books have been a constant source of inspiration and a model of clarity. Despite our substantial disagreement over important issues, he has always helped me to make my arguments as forceful as possible, even when those consisted in direct criticisms of his own view.

Will McNeill read large portions of a late draft of the thesis. He greatly helped me to improve the structure of the work at important junctures, and he saved me from a number of mistakes.

Michael Sollberger has read almost everything that I have written during my doctoral studies. He has always been patient, generous and very open to my philosophical ideas, even when these were formulated in a very confusing way in some of the early drafts.

Jorgen Dyrstad read a late draft of my work with astonishing rapidity, and returned it with an incredible amount of comments. His comprehensive knowledge and understanding of the literature have been extremely valuable to me, and his enthusiasm for philosophy is highly contagious.

I would also like to thank Davor Bodrozic, Filipe Carijo, Tim Crane, Coralie Dorsaz, Fabian Dorsch, Giulia Felappi, Ivan Ivanov, Dave Jenkins, Fiona Macpherson, Martine Nida-Rümelin and Mattia Sorgon.

Writing this thesis was made possible by the Swiss National Science Foundation (research grants Nos. 100018_129598 and P1FRP1_148549).

Abstract

When we perceive we are under the impression of being directly aware of concrete, mind-independent objects. We also consider perception as a basic, reliable source for acquiring beliefs and an effective means for coping with the environment. In the philosophical literature, this direct and basic character of perception is sometimes captured by saying that perception is *openness to the world*. Articulating, refining and vindicating as far as possible this commonsensical view of perception as openness to the world is the main objective of this work.

In order to make the metaphor of openness tractable in theoretical terms, I set up the dialectic in terms of a contrast between representationalism, which holds that perception represents things as being a certain way, and pure relationalism, which holds that perception can be openness to the world only by being a non-representational relation between a subject and objects in the world. The main thesis that I defend with several arguments in this work is that perception is both representational and relational. Indeed, it intelligibly relates us to the world in virtue of its representational content.

In chapter one, I place perception in the broader context of intentional states. I articulate a view on which intentional states are directed toward their objects in virtue of having representational content, and show that *prima facie* the arguments for thinking that paradigmatic intentional states such as belief have content apply to perception too. I also explain why pure relationalists would reject the application of the content-model to perception.

Chapter two is devoted to Charles Travis's argument that perceptual experience does not have content. I counter his argument by offering a phenomenological account of perceptual constancy, and I show that content is necessary for capturing constancy in theoretical terms. Further, I argue that content helps us to better explain how the mind-independent nature of concrete objects is manifest in the phenomenology of perceptual experience.

Chapter three concerns the question of how we should elucidate perceptual experience as openness to the world in the face of illusion and hallucination. I consider three possible approaches (sense-data theories, representationalism, and pure relationalism), and I argue that a view on which perceptual experience is both representational and relational can best make sense of the relevant phenomena.

In the fourth chapter, I consider the debate on phenomenal character between representationalist and purely relationalist views in light of a distinction between particularity and generality. I show how representationalists are committed to there being an element of generality in perception, and how pure relationalists construe perception as wholly particular instead. I defend the representationalist view on the grounds that generality is necessary to make sense of the phenomenal and epistemic determinacy of perception.

In the concluding remarks, I consider the question of whether my view can make room for the relationality of perception, and I argue again that only a distorted view of the metaphor of openness to the world could lead one to think that content and relationality are incompatible.

Introduction

This is a work on the intentionality of perception. The topic of intentionality is how mental states can be about the world, and thus the topic of this thesis is how our perceptual experiences can be about the world and the things it contains. The question lends itself to different interpretations, so I should take this opportunity to clarify that I intend to discuss the intentionality of perception in a particular way.

There is one influential approach to the question that I want to mention just to set it aside. Over the last few decades, a number of research programs have interpreted the problem of intentionality as that of reducing intentional notions such as reference, content and truth to non-intentional notions more in tune with natural sciences, such as causality, information and (a specific understanding of) representation. The aim of these research programs, pioneered by Fred Dretske (1981), Jerry Fodor (1987) and Ruth Millikan (1984), has been to naturalize intentionality, that is, to conceive of intentionality in such a way as to render it compatible with a scientific view of the world and human beings. Although these programs have been exciting and fruitful, the naturalization of the intentionality of perception is *not* my primary concern here. Rather, I aim to elucidate perception as *openness to the world*.

Openness to the world is a metaphor that has been popularized in the analytic world by John McDowell (1994), which is meant to capture the direct and basic character of our perceptual encounter with the world: when we perceive, there is nothing standing between the world and us, and perceiving the world is the most basic source for our conception of concrete reality. It has to be acknowledged that both directness and basicness could be captured within a naturalistic framework. For example, one could capture the fact that perception is direct by emphasizing that perception has the function of conveying information about concrete objects, rather than, say, light arrays. One could also capture the basicness of perception by saying that the intentionality of perception is, unlike that of artifacts and language, *original*, or *undervived* – that is, its intentionality does not depend on the agent's possession of intentional capacities.¹ Why do I prefer a metaphor to scientific language?

The difference between a naturalistic approach and what I do in this thesis could be characterized as the difference between an explanatory task and a descriptive task. The naturalization of intentionality is an explanatory project, in the sense that it aims to explain a

¹ This is Dretske's view (1981, 1995).

set of intentional notions in terms of other, non-intentional notions. I rather adopt a descriptive approach, in the sense that I aim to elucidate perception from a subjective point of view, taking at face value and systematizing our ordinary, pre-theoretical conception of perceptual experience, without trying to reduce the concepts that inform the pre-theoretical view to anything more third-personal or objective. Although it could be disputed whether there is anything like a unitary, commonsensical conception of perceptual experience and its objects, it seems that propositions such as that things look differently when viewed from different angles, or that sometimes things are not as they look to be, or that we have a sense of perceiving things that are there independently from our perceiving them, would all be accepted upon reflection by any adult human being – this suggests that there is a pre-theoretical view of experience which we implicitly accept in our dealings with the world. Refining such ordinary conception and offering a theoretical gloss on it is the way in which I address the intentionality of perception. In this connection, retaining and elaborating on the metaphor of openness to the world is a way of anchoring the theoretical apparatus that I employ throughout the thesis to a more basic set of notions that are subjectively salient, in the sense that a subject would be able to articulate them upon reflection, without assuming any particular theoretical outlook on perception or employing observational evidence coming from natural sciences.

In slightly different words, it could be said that this work is an exercise in phenomenology rather than philosophy of science. The basics of the phenomenological method have been effectively summarized by Charles Siewert, who has identified four conditions that you have to satisfy if you want to do phenomenology:

1. You explain mental or psychological distinctions.
2. You show how such explanation has significant theoretical consequences.
3. Your explanation relies on a source of warrant special to some first-person applications of the distinctions explained.
4. You do not assume such first-person warrant as you rely on is derived from observational, third-person evidence. (Siewert 2007: 202)

On a plausible interpretation, an example of the kind of mental distinctions that Siewert has in mind could be the distinction between *what* one perceives on the one hand, and the *way* in which one perceives it on the other: this distinction concerns mentality, it has theoretical

consequences (or so I argue, especially in chapter two), and it relies on a direct, first-personal access that one has to one's experience.²

Although I aim at articulating and vindicating as far as possible an ordinary, commonsensical view of perception, the kind of elucidation that I offer is theoretical, so I make heavy use of a theoretical apparatus that is well-established among contemporary philosophers of perception. A relevant theoretical contrast to be drawn at this point is that between those who think that perception is fundamentally representational, those who think that it is fundamentally relational, and those who think that it is fundamentally both representational and relational, as I do.

According to *representationalism*, or *intentionalism*, or the *content view*³, perception has representational content; that is, it represents the world as being a certain way. Contents are normally taken to be the sort of thing that can be assessed for truth or accuracy, and thus can make an experience veridical/falsidical, or accurate/inaccurate⁴. In other words, an experience represents accurately or inaccurately in virtue of its content:

On such a view, perceptual states represent to the subject how her environment and body are. The content of perceptual experiences is how the world is represented to be. Perceptual experiences are then counted as illusory or veridical depending on whether the content is correct and the world is as represented. (Martin 1994: 464)

Here too, I should emphasize that although one major motivation for putting forward a representational account of perception is that representation could be analyzed in terms of the naturalistic notions of causality and information, this is not the main reason why I discuss representationalism and ultimately endorse a version of it. Rather, I am interested in representational content as a useful theoretical tool for explicating perception as openness to

² Siewert uses the word “explanation”, but it is clear that by that he means something like “conducive to understanding”, not the stronger idea of reducing certain notions to others.

³ These terms are often used as being different from one another in meaning, but for present purposes they can be treated as synonyms, as the representationalist, the intentionalist and the content view theorist all claim that perception has representational content – this minimal claim is the one that interests me.

⁴ Sometimes it is claimed that perceptual experiences are better characterized as being accurate or inaccurate, rather than veridical or falsidical, because an experience can represent things *more or less* accurately, whereas truth and falsity do not come in degrees. Nothing in my arguments hinges on that, but the interested reader can consult Crane (2009) for a detailed discussion.

the world and capturing the kind of distinctions that Siewert thinks a phenomenologist should draw.

The idea that perception can represent or misrepresent has enjoyed a large consensus at last since the eighties. However, an alternative view that has recently gained significant popularity holds that perception is relational *instead* of being representational. It holds, that is, that perceptual experience is fundamentally a question of a subject being in a perceptual, non-representational relation to mind-independent objects; let us call this view *pure relationalism*, in order to distinguish it from a view on which the relationality of perception is compatible with its being representational.⁵ According to pure relationalism, perception makes contentful representation of reality possible. However, it does so precisely in virtue of being more basic than representation. Thus Brewer writes:

Perceiving is not a matter of being saddled with representational content...It is rather a matter of the conscious presentation of actual constituents of physical reality themselves, particular such things, just as they are, which is what makes all contentful representation of that reality in thought even so much as possible. (Brewer 2006: 172)

My exploration of the intentionality of perception is to a large extent informed by the contrast between representationalism and pure relationalism. More in detail, the structure of the thesis is the following: in chapter one, I expound a view of intentionality that is by far and large uncontroversial, which essentially involves the notion of content. Then I explain why the considerations that count in favour of introducing content for paradigmatic intentional states such as belief or desire do seem to apply to perception as well. I also explain how purely relationalist authors such as Brewer and Travis reject the application of the content-model of intentionality to perception, and thus I set up the dialectical contrast between representationalism and pure relationalism to be discussed in the subsequent chapters.

Chapter two is devoted to the work of Charles Travis, and especially to his argument that perceptual experience does not have content. After a detailed discussion of the argument, I offer a phenomenological account of perceptual constancy, and I claim that we can only make sense of constancy at a theoretical level if we rely on the notion of content. Further, I argue

⁵ For some examples of representationalist views see, Burge (2010), Byrne (2001, 2009), Crane (2001, 2006, 2009), Dretske (1981, 1995), Searle (1983) and Siegel (2010). For purely relationalist views see Brewer (2006, 2008, 2011, 2015), Campbell (2002a, 2002b), Fish (2009) and Travis (2004, 2007, 2008, 2013). Logue (2014), McDowell (2013) and Schellenberg (2011) emphasize that perception is both representational and relational.

that content helps us to better explain the way in which the mind-independent nature of concrete objects is manifest in the phenomenology of perceptual experience.

Chapter three concerns the question of how one should elucidate perceptual experience as openness to the world in the face of illusion and hallucination. I consider three possible approaches (sense-data theories, representationalism, and pure relationalism), and I argue that a view on which perceptual experience is both representational and relational can best make sense of the relevant phenomena. I also discuss three purely relationalist accounts of illusion, and I argue that Brewer's sophisticated theory concerning the way(s) objects look is the best chance for those who want to dispense with content, although it is not exempt from problems.

In the fourth chapter, I offer a distinctive interpretation of the way in which representationalist and purely relationalist views disagree over the topic of phenomenal character, by considering the debate in light of a distinction between particularity and generality. I discuss an argument for content offered by Susanna Siegel, and then I go on to criticize it as being question begging. I then argue that what really underlies her argument is a commitment to there being an element of generality in perception, and I show how pure relationalists would resist the idea by construing perception as wholly particular. Finally, I argue that generality is indeed necessary to make sense of the phenomenal and epistemic determinacy of perception, thus indirectly rehabilitating Siegel's argument.

In the conclusion, I consider again the question whether my view can make room for the relationality of perception, and again I argue that only a distorted view of the metaphor of openness to the world could lead one to think that content and relationality are incompatible. Perception is in fact fundamentally both relational and representational, and it discloses the world to our view precisely in virtue of its content.

Chapter One: **Intentionality**

Abstract

Intentionality is the property that mental states have of being directed toward something. The objective of this chapter is to develop a particular version of this claim, on which mental states are so directed in virtue of having representational content. In the first section, I articulate the notion of intentional object, and in the second section that of intentional content. In the third, I connect content to the notions of conditions of satisfaction and intensionality-with-an-s. In the fourth and final section, I apply to perception the conception of intentionality that emerges from the earlier sections of the chapter. There, I also explain how pure relationalists such as Bill Brewer and Charles Travis would resist the application to perception of the content-model, thus setting up the dialectical contrast between representationalist and purely relationalist views of perception to be investigated in the rest of the thesis.

Introduction

Every discussion of intentionality that I'm aware of begins with the observation that the mind has the capacity of being *directed* toward the world, or that subjects have mental states which are *about* the world or objects in the world – call this capacity *direction*. Coupled with this, goes the observation that the mind can take different *attitudes* toward its objects, where paradigmatic examples of such attitudes are love, fear, desire, admiration, thought, judgment, and belief – call this capacity *attitude*. *Direction* and *attitude* are the basic elements of intentionality, and that much seems agreed on all sides. The *locus classicus* for the formulation of these points is to be found in Brentano's *Psychology from an Empirical Standpoint*:

Every mental phenomenon is characterised by what the Scholastics of the Middle Ages called the intentional (or mental) inexistence of an object, and what we might call, though not wholly unambiguously, reference to a content, direction towards an object (which is not to be understood here as meaning a thing) or immanent objectivity. Every mental phenomenon includes something as object within itself, although they do not all do so in the same way. In presentation, something is presented, in judgement something is

affirmed or denied, in love loved, in hate hated, in desire desired and so on. (Brentano 1995: 88)

The passage is, of course, fraught with exegetical danger due to the presence of expressions like “intentional inexistence”, “reference to a content”, and “immanent objectivity”.¹ However, *direction* and *attitude* emerge quite clearly when Brentano writes of a “direction towards an object”, and when he remarks that in presentation something is presented, in judgement something is affirmed or denied, and so on. Contemporary discussions of intentionality take *direction* and *attitude* as a starting point too. Thus Searle comments on both elements:

Intentionality is that property of many mental states and events by which they are directed at or about or of objects and states of affairs in the world. If, for example, I have a belief, it must be a belief that such and such is the case; if I have a fear, it must be a fear of something or that something will occur; if I have a desire, it must be a desire to do something or that something should happen or be the case; if I have an intention, it must be an intention to do something. (Searle 1983: 1)

And thus Zahavi expounds Husserl’s theory of intentionality:

In his analysis of the structure of experience, Husserl pays particular attention to a group of experiences that are all characterized by being conscious of something, that is, which all possess an *object-directedness*. This attribute is also called *intentionality*. One does not merely love, fear, see, or judge, one loves a beloved, fears something fearful, *sees* an object, and judges a state of affairs. (Zahavi 2003: 14)

There are many questions related to the two basic elements of intentionality that I have identified, so I will be highly selective in my choice of topics to discuss. My aim is not so much to provide an overview of all the approaches to intentionality, as to articulate a general conception of intentionality, as uncontroversial as possible (although not entirely so), that will provide a framework for a more specific exploration of perceptual intentionality throughout the rest of this work. In the first section of this chapter, I specify *direction* through an

¹ It is worth noting that Brentano thought that intentionality was the mark of the mental, and thus that *all* (and only) mental states exhibit intentionality. While this is seconded by some contemporary philosophers, such as Crane (2001), there are others, such as Searle (1983), who think that only some mental states exhibit intentionality. None of my arguments hinge on this.

explication of the notion of intentional object, and I give a *prima facie* criterion for the individuation of intentional states. In the second section, I argue that in fact the notions of intentional attitude and intentional object are insufficient to individuate intentional states correctly, and I give two arguments for the claim that a third, additional notion is needed: *representational content*. More specifically, I emphasize that content helps us to capture in theoretical terms a pervasive phenomenon that anyone could recognize upon reflection on one's own mental life: when we are directed toward an object, our so being directed typically implies a certain *perspective* on it, where this can be understood either in literal terms, as in the case of perception or visualization, or in metaphorical terms in the case of other intentional states. In the third section, I expand on the idea of content and relate it to the notions of *conditions of satisfaction* and *intensionality-with-an-s*. In the fourth section, I turn to perception and apply to it the notion of intentionality previously articulated.

1. Objects

The aim of a theory of intentionality should be that of understanding what it means for our language, thought and perception to be directed on or be about the world. On one interpretation, this is a question about reference: it asks what it means for an expression or a perception to refer, and thus one should begin a theory of intentionality with a theory of reference, and discuss things such as perceptual reference, demonstrative reference, singular and general reference, proper names and descriptions. While this certainly appears a very sensible method for explicating fundamental forms of intentionality, and will be relevant at many points in this work to the characterization of the intentionality of perception, it also seems that the theme of intentionality in its most general form is broader than reference, in that reference is a relation, and as such implies the existence of the *relata*, whereas *being about something* is the real theme of intentionality, and *being about something* is a predicate that can be satisfied by states directed at objects of pure phantasy, as when one thinks about Superman, or is afraid of the monster under the bed – that is, by episodes that fail to refer, but are as intentional as anything can be.

What is it for a state to be about something then, if it is not necessarily to refer? There are two complementary ways of explicating what it means to say that mental states have a direction toward, or are about, something. One is to elaborate on the nature of the objects that

our mental states purportedly are about, and another is to elaborate on what it is to be intentionally related to such objects. I start with the first, for illuminating the nature of the relatum will prove useful in understanding the nature of the relation, although it will emerge that in *some* cases “relation” might be a misleading term.

I assume that every intentional state has an (intentional) object. This seems plausible to me, as it pertains to the very idea of intentionality that in having intentional states one is directed toward something. What does “intentional object” mean, though? One thing should be made clear right from the start: there is no special category of intentional objects, either in addition or in opposition to other kinds of objects such as, say, abstract objects or concrete objects. Thus Crane writes:

Intentional objects are, by definition, those things in the world which we think about; or those things which we take, or pretend, or otherwise represent to be in the world; or which we merely represent in thought. If there is such a thing as thinking about “things” (in these various senses of “thing”) then there are intentional objects. (Crane 2013: 4)

The characterization offered by Crane is completely neutral with respect to the question of what kinds of ontological features intentional objects have, and this for the good reason that there is no one unifying feature of all intentional objects. Speaking of intentional objects is perfectly compatible with the reasonable, realist insight that when we think about and refer to objects that actually are in our surroundings, it is those very objects, and not some special intentional object, that are the objects of our attitudes. Indeed, in a swathe of central cases, intentional objects are ordinary, concrete objects such as tables, chairs, and persons. The need for accommodating this basic, pre-theoretical point about the way we are directed toward things leads Searle to simply equating intentional objects with ordinary objects:

...an Intentional object is just *an object like any other*; it has no peculiar ontological status at all. To call something an Intentional object is just to say that it is what some Intentional state is about. Thus, for example, if Bill admires President Carter, then the Intentional object of his admiration is President Carter, the actual man and not some shadowy intermediate entity between Bill and the man. (Searle 1983: 16-17, emphasis mine)

Avoiding shadowy intermediate entities can only be right, and equating intentional objects with ordinary objects is certainly apt to characterize many paradigmatic cases of intentionality. However, if we want to give the most general possible characterization of intentionality, saying that an intentional object is *an object like any other* is not completely exempt from complications. For one thing, if *an object like any other* means *any object that exists*, then this leaves unexplained how we can be directed toward objects that do not exist. Relatedly, sometimes intentional objects are indeterminate, whereas ordinary objects are not: for example, I can think or imagine that there is a man around the corner without thinking or imagining a man of a particular height, whereas I could not hit a man without hitting a man of a particular height. Finally, if an intentional object is an object like any other in the sense that intentional objects are concrete objects, this seems far too restrictive, because we can also be directed toward properties, states of affairs, events, fictional entities, and all sort of abstract objects: if I believe that there's a red ball, my belief has as its object – is directed on, concerns, is about – a red ball. If I hope that the party will be nice, the party is the object of my hope. If I fear that the theory of relativity is too difficult for me, the theory of relativity is the object of my fear. If I admire Superman, then Superman is the object of my admiration.

Crane (2001) has argued that intentional objects are to be understood in terms of a *schematic conception*, as opposed to a *substantial conception* of objects. A substantial conception of a kind of object aims at giving a unitary characterization of the kind in question. In articulating a substantial conception of, say, mathematical objects, one might say that having a spatiotemporal dimension is not something which pertains to their nature (the number two surely doesn't have any spatiotemporal location, as opposed to a material object like a red ball). Not being located in space and time, then, might constitute one defining feature of all mathematical objects.

The articulation of a schematic conception of a kind of object has a more modest aim, for it simply is an attempt at explaining why we group certain objects under a certain heading – *intentional objects*, say – despite the lack of any unifying feature among those objects.²

Crane explicates the notion of intentional object, in the sense of schematic objects, by means of an analogy with *objects of attention*. Just as the correlates of our attention vary in

² One might say that their being objects of our mind is a unifying feature. However, this does not contribute to the individuation of an ontological kind any more than a feature such as *being to the left of John* does – that is, an object may or may not be an object of our mind on one occasion or another, without that affecting in any way the kind of object that it is; *being an object of the mind* is not a feature that can be used in order to carve nature at its joints, so to speak.

kind from one occasion to the other, so do intentional objects. In the case of intentional objects, we group together objects belonging to very different ontological categories by means of the label “intentional objects”, insofar as they all can be what the mind is directed upon in one circumstance or another – equivalently, they are whatever one would name in answering the question: “what is your belief (hope, desire...) about?”.³

I acknowledge that the notion of schematic object is not maximally clear. However, all what I need for my purposes is to get the reader to agree on the phenomena in question, namely that intentional objects are heterogeneous, and that a thought about Superman is about an object in the *same* sense in which Bill’s admiration for Carter is about Carter – “about” does not change its meaning from one case to the other.⁴ If someone had a better explanation than schematic objects for the phenomena in question, I would be happy to accept it.

It seems that any intentional state should be individuated at least in part with reference to the object that the state purports to be about, for that is the whole point of *direction* – our mental states concern objects, in the schematic sense specified above. But the mind also has the capacity to bear different *intentional attitudes*⁵ to the very same intentional object. I might *believe* that Mary is beautiful, *hope* that Mary is single, *desire* that Mary and I go out together, and so on. Obviously the converse also holds: I *believe* that the ball is red, I *believe* that the number four is even, and so on. It is intuitive to take the object/attitude pair as a *prima facie* criterion for the individuation of intentional states: for any two intentional states, we can say that they are instances of the same type *iff* i) they have the same object and ii) the object is intended in the same attitude. Thus you and I would be in the same (type of) intentional state

³ Another analogy employed by Crane (2001:15) is that between intentional objects and grammatical objects. As in the case of intentional objects, here too we find an unrestricted variety of kinds of things that can be correctly thought of as grammatical objects, despite the lack of any unifying feature. The analogy is not without problems though, since a grammatical object is *a word*, and that word exists in all cases, although it may or may not refer to various objects of various sorts – but then all the explanatory burden is carried by the notion of *reference*, rather than by that of intentional object.

⁴ One might object to this as follows: in cases of reference to an object, we can make sense of thinking *true* thoughts about the object in question, and the truth of those thoughts is grounded in the way the object is. However, the objection goes, a non-existent object cannot ground any truth, precisely because it does not exist. Therefore, intentional states cannot be about non-existent objects in the same sense in which they are about objects of reference. Addressing the issue of truths about non-existents would require an extensive discussion of questions that do not fall within the scope of this chapter. However, it is worth pointing out that responding to the objection does not necessarily require one to embrace a controversial Meinongian ontology (although see Berto (2013) for an articulated defence of it). Indeed, Tim Crane (2013, chapter 5) has argued in favour of a reductive approach to truths about non-existents – that is, for the idea that truths about non-existents should in principle be reducible to truths about existents, such as truths about *representations* of non-existent objects.

⁵ These are also called *intentional modes* (Searle 1983), but I use the term “attitude” to avoid any terminological confusion with *modes of presentation*, which I introduce later.

if we both desired that Mary was single, whereas we would be in different intentional states in case I *desired* that Mary was single and you *believed* her to be so. In fact, this is a very rough guide to the conditions of identity of intentional states, for it is possible for two subjects to be in two different intentional states despite sameness of object and attitude. The additional ingredient that is needed for a characterization of intentionality is *content*.

2. Content

There are at least two classic arguments for introducing the notion of content: one has to do with the fact that some intentional objects do not exist — that is only marginally relevant for our purposes, and thus I examine it quite briefly. The other, which is more central to this thesis and will be developed to greater length in chapter two, has to do with the fact that intentional states with the same object and attitude can still differ in significant epistemic and phenomenal respects.

2.1 *Non-Existent Objects*

Suppose that one tried to individuate intentional states solely in terms of the attitude/object pair. Walter Hopp (2011) labels this view the “spotlight view” (13), for an attitude would be like a spotlight that falls on one object or another, and the spotlight/object pair would be sufficient to distinguish the state from any other possible (type of) intentional state. However, as Hopp points out, it’s not difficult to see how states whose intentional objects do not exist cause trouble for this view. The spotlight view has to account for *direction* – for how a mental state can be about something – simply by appealing to an attitude and an object. But if the object doesn’t exist, it’s not clear how it can play an explanatory role with respect the question of the state’s aboutness.

More precisely, the challenge for the spotlight view is to make sense of the *differences* among intentional states with objects that don’t exist. For even if one were to concede that the spotlight view can account for mental states with nonexistent objects – the image would be that of “a spotlight shining into a void” (Hopp 2011: 18) – it would be hard to see how it could account for mental states with *different* nonexistent objects; the spotlight would always shine into the same void, as it were. Mental states with different nonexistent objects have

different inferential, behavioral, and emotional implications. If I admire Superman, for example, I will buy all the Superman comic books, get excited about the new Superman movie, and so on. If I admire Batman, I will buy all the Batman comic books, get excited about the new Batman movie, and so on. Generally speaking, a certain attitude toward an object can rationalize or imply certain other attitudes, but it's difficult to see how a void could rationalize or imply anything, let alone how the same void could rationalize or imply distinct sets of attitudes.

It will not work to fill the void with an object in the mind, thus conceiving intentional states concerning objects that don't exist outside as relations to objects that exist "inside", whatever that meant. This immanentist conception of intentional objects, on which intentional states *always* incorporate a reference to *mental objects*, is often attributed to Brentano, and was already convincingly criticized by Husserl (2001: 97-100; 125-127). The immanentist conception suffers from obvious ontological, phenomenological and epistemological problems. For one thing, the very notion of mental object is very elusive. For another, if Timmy hopes that "Rudolph the Red-Nosed Reindeer lands on his roof tonight", something Tim is definitely *not* hoping for is that "one of his ideas or an abstract object will land on his roof tonight" (Hopp 2011: 14), so this would get the phenomenology of Tim's hope entirely wrong. With respect to epistemological concerns, it is not clear how a relation to a mental object could explain our knowledge of mind-independent objects, and modelling the relation between mental objects and mind-independent objects in terms of some similarity between the two is notoriously problematic.

Thoughts about non-existent objects cannot be *relations* between subjects and objects, because the obtaining of a relation entails the existence of its *relata*, and non-existent objects do not exist, either mentally or extra-mentally. I think that the right category to characterize thoughts about non-existent objects is that of representation. That in virtue of which an intentional state can be about a non-existent object, in the same sense in which admiring Carter is about Carter, is that the state represents (by means of its content) that there is an (non-existent) object which is a certain way. This facilitates the task of making sense of how we can be directed toward objects that don't exist, because, generally speaking, if X represents Y, then Y need not exist.

In the next subsection, I explain representational content in terms of *Fregean senses*, or *modes of presentation*, and I argue that *all* intentional states (not only those directed at non-existent objects) have representational content.

2.2 Modes of Presentation

Intuitively, if a subject is directed on the world, then the subject has a *perspective* on the world. Crane highlights that there are two features of perspective, understood in a somewhat narrow, pictorial sense, that carry over to a broader use of the term, on which having a perspective on the world includes things like having an opinion, a way of seeing things, a stand on certain matters, and so on. One feature is that a perspective is “*on things* other than the perspective itself” (Crane 2001: 6), where this is mirrored – in the broader sense in which intentionality is a perspective on the world – by the fact that intentional states are about, or directed on, something (typically something other than themselves). Second, perspective is such that things in a picture are necessarily “presented in a *certain way*” (6); it belongs to the idea of perspective that there are certain things, or aspects of things, which are included in the picture, and others that are excluded. This is mirrored by the fact that there are different ways of being directed on the same object. It is this second feature which is especially relevant to the notion of content.

For any two intentional states in the same attitude and with the same object, there can still be significant epistemic and phenomenological differences between them if the states involve a different perspective – in a suitably broad sense – on the object. The notion of perspective at stake can be sharpened by means of the notion of a Fregean sense. If I think about the morning star, while you think about the evening star, our states share the object (because the morning star is identical with the evening star) and the attitude (because they are both thoughts). And yet, the way in which they intend the object is different: the object is intended in the two states according to two different modes of presentation, or equivalently, according to two Fregean senses. A Fregean sense, or mode of presentation, is a way of referring or purporting to refer to an object. Objects, in the broad sense of intentional objects, have different aspects that can be captured in the various ways that we have of thinking about them. An aspect of the morning star surely is that it appears in the morning, whereas it is an aspect of the evening star that it appears in the evening. Since the two stars (let us pretend we

don't know that in fact it's the planet Venus) are in fact the same, it follows that the two aspects belong to the very same object. These two different aspects of the object yield two different ways of thinking about it: the aspect *appearing in the morning* might yield thoughts about the object conceived as the morning star, while the aspect *appearing in the evening* might yield thoughts about the object conceived as the evening star; these two thoughts are surely different in epistemic (and perhaps phenomenological) respects.⁶

Generally speaking, aspects are properties of objects, such as having a pointed nose, or appearing in the morning, or being disposed to behave a certain way, or having a certain political function, as when a city is the Capital of a State or a man is the Mayor of that city. They can also be incorrectly ascribed to objects: I might think of a certain man as the boss of the local criminal gang even if he isn't.

It is useful to think of content in terms of aspects and modes of presentations: when we are intentionally directed on an object, and a certain aspect of that object is the epistemic or phenomenological focus of our so being directed, we can say that the object is intended according to a mode of presentation that corresponds to the aspect in question. Thus, when we think about a star under the aspect that it has of appearing in the morning, this makes it so that the star is thought of according to the mode of presentation *the morning star*, and that mode of presentation constitutes the intentional content of the state⁷.

Fregean senses or modes of presentation pick out objects under certain aspects. This means that they identify an object by specifying some of its properties, and thus provide a *description* of the object. On one understanding, this means that they refer satisfactionally, that is, they refer to whatever object satisfies the description. However, this picture of how we perceive and think about objects is problematic, because it fails to accommodate that we think about, perceive or otherwise refer to *particular* objects, and not just to whatever object has

⁶ I do not need to take a stand here on whether there is a phenomenology of thought, but phenomenological differences between different ways of being directed toward the same object will become crucial when discussing perceptual content in chapter two.

⁷ As in the case of intentional objects conceived as schematic objects, here too I do not want to rely too much on a specific theoretical framework, but rather to get the reader to agree on a phenomenon, namely that there are significantly different ways of being directed toward an object. An important alternative to Fregean contents are Russellian propositions, whose constituents are individuals, properties and relations. Crucially though, those who endorse Russellian contents too have to make sense of the phenomenon which I capture in terms of Fregean contents, and they typically do so by means of the notion of *guise* – for them, for example, a belief is a ternary relation “among believers, propositions, and some third type of entity”...“perhaps something like proposition guises, or modes of acquaintance or familiarity with a proposition, or ways in which a believer may take a given proposition” (Salmon 1989: 246).

certain properties. This was effectively pointed out by Strawson (1959) by means of the *reduplication argument*, which has then been extensively commented on by Brewer (1999), among others.⁸ Brewer invites us to consider a complex description, such as “The [unique] red ball under the glass table between the chair and the sofa in front of the round white window”. Now, an interesting way in which the description would fail to identify a particular mind-independent object would occur if there was more than one object that satisfied the description, that is, if there was somewhere in the universe, a reduplication of the scene described:

...however detailed and extensive the description may be, there is always the possibility, in principle at least, of a reduplication elsewhere in the universe of the whole scene as described...So the belief [that there is a red ball under the glass table] fails altogether to be a belief that any *particular* mind-independent thing is determinately thus and so. (Brewer 1999: 27, emphasis mine)

Does that mean that Fregean senses are unfit to constitute contents? This would be too quick a conclusion. Modes of presentation are a formidable tool for modelling intentionality, because they allow us to capture the perspectival dimension of intentionality. But it is certainly not extraneous to the idea of a perspective on things, either in the narrow pictorial sense or in the metaphorical use that is relevant to present purposes, that the perspective is a view on particular things. So we have to find a way of keeping together two elements of perspective: the focus on a particular object, and the general cognitive outlook that comes with a description of that object – general because things other than the particular object in question might satisfy the description.

Sometimes my friends refer to me as “the guy who always closes the window because he suffers from sinusitis”. While almost certainly I am not the only person in the universe who closes windows and has sinusitis, it seems obvious that my friends refer to me, and not to whoever satisfies the description; they take a particular perspective on *me* and not someone else. In order for them to be able to do that, something in the content of their utterance must anchor the description to a particular individual. McDowell (1984) has argued, explicitly relying on Evans (1982), that it is perfectly possible to entertain *de re senses*, that is, senses

⁸ For lines of thought similar to Brewer’s, see Evans (1982) and Eilan (1997).

that are specific to their *res*. What makes them so specific is, according to McDowell, a *context*:

...if a propositional attitude (or utterance) is essentially *de re*, that is in virtue of the fact that a context involving the *res* enters itself into determining how the attitude (or utterance) can be correctly ascribed. (McDowell 1984: 284, footnote omitted)

Why is it appropriate to attribute to my friends attitudes about me, and not about any person with sinusitis? Because a context in which I figured anchored their attitudes to a particular individual (that is, me), rather than to anyone who satisfied the relevant description. Naturally, it is crucial that their relation to the context and to myself be not understood in *entirely* descriptive terms, on pain of the predicament described by Strawson and Brewer as *reduplication*. Suppose, for example, that one tried to explain why my friends refer to me by enriching the description along the following lines: the guy with sinusitis who always closes the window and sits by the brown table and wears a red pullover and black glasses. Just as in the case of the red ball under the glass table, here too there is nothing that could possibly rule out that someone else in the universe satisfies this very same description. The moral is that my friends' relation to me and the context in which we are all embedded has to be understood in genuinely relational terms, where this involves at the most basic level a demonstrative identification of myself and other elements to which I stand in spatial relation. Thus, what ultimately determines that my friends refer to me by means of the description "the guy who always closes the window because he suffers from sinusitis", either when actually seeing me closing the window or simply talking about me on a different occasion, is the fact that they stand or stood in a perceptual relation to me and other elements in the environment, and that relation provides or provided an informational link sufficient to single me out as the referent of the general description.

We are now in a position to refine our criterion for the individuation of intentional states. For any two intentional states, they are the same just in case the following three conditions are met: i) they have the same content; ii) they have the same object⁹; iii) the object is intended in the same attitude.

⁹ In paradigmatic cases where we refer to objects in our environment the conditions of sameness of object becomes redundant, since I conceive Fregean senses along the lines suggested by McDowell (1984).

In the next section, I elaborate on the idea of content, with reference both to the notion of conditions of satisfaction of a state and to the notion of intensionality-with-an-s. This will enable us to better understand the perspectival nature of intentionality and its relation to configurations of the world.

3. Satisfaction Conditions, Intentionality and Intensionality-with-an-s

An intentional state includes a particular way of conceiving things, where things actually may or may not be that way. If I believe that there's a brown table in my room, then *according* to my belief state there's a brown table out there, where in fact it might turn out not to be one. Generalising, we can say that intentional states have conditions of satisfaction: the content of a state determines its conditions of satisfaction, and these conditions are what has to be the case if the state is to be satisfied, or correct. Thus specifying the content is a way of specifying the conditions of satisfaction. A state is satisfied, or accurate, *iff* things actually are the way they are intended to be in the particular perspectival take on things which is included in the state. Thus Searle says:

My belief will be satisfied if and only if things are as I believe them to be, my desires will be satisfied if and only if they are fulfilled, my intentions will be satisfied if and only if they are carried out. (Searle 1983: 10)

The relation between content and conditions of satisfaction has to be handled with care though, for not any specification of the satisfaction conditions will be equally good as a specification of the content. This brings me to the connection between intentionality and intensionality-with-an-s.

Intensionality-with-an-s is a property of some linguistic contexts, and its presence can be ascertained by means of simple tests: if a certain linguistic context doesn't license existential generalization or substitution of co-referring terms *salva veritate*, then it is intensional-with-an-s.

The rule of existential generalization says that from:

(1) F(a)

It can be inferred:

(2) $\exists x F(x)$

For example, from:

(3) John is tall

It can be inferred:

(4) Someone is tall

But that doesn't hold true of intensional contexts, examples of which are generated by some of the verbs used in attributing intentional states, like "believing" or "desiring". For example, from:

(5) John believes that Superman flies

It cannot be inferred:

(6) There is some x such that John believes that x flies

The principle of substitution of co-referring terms says that from:

(7) $F(a)$

(8) $a = b$

It can be inferred:

(9) $F(b)$

For example, from:

(10) Superman flies

(11) Superman = Clark Kent

It can be inferred:

(12) Clark Kent flies

But this principle too fails in intensional contexts, for from:

(13) John believes that Superman flies

(14) Superman = Clark Kent

It cannot be inferred:

(15) John believes that Clark Kent flies

It is well-established that intentionality and intensionality-with-an-s are two different phenomena. For one thing, there are intensional contexts that have nothing to do with intentionality. For example, modal operators generate intensional contexts. From:

(16) The number of coins in my pocket is five

(17) Five is necessarily odd

It cannot be inferred:

(18) The number of coins in my pocket is necessarily odd (Crane 2001: 12)

For another, not all reports of intentional states generate intensional contexts, for if Lois Lane loves Superman, then Lois Lane loves Clark Kent.

However, it would be too drastic to assert that there are no interesting connections at all between intentionality and intensionality-with-an-s. In fact, I believe that both the failure of existential generalization and that of substitutivity are logico-linguistic reflections of features that intentionality actually has: failure of existential generalization reflects the fact that intentional objects need not exist, and failure of substitutivity reflects the fact that intentionality is perspectival.

The perspectival nature of intentionality and the failure of substitutivity together illuminate why not any specification of the satisfaction conditions of an intentional state will do a good job in expressing its content. Satisfaction conditions are a question of what is the case, and as such they are neither extensional nor intensional. However, extensional reports of, say, belief, might often fail to reflect the perspectival nature of belief that contributes to making it *intentional*. This is not to say that extensional reports are illegitimate, for it would certainly be permissible to say, of John, the following:

(19) Of that man, Clark Kent: John believes he can fly

In this extensional, *de re* attribution of belief John is put in relation to an individual (rather than a whole proposition), where the individual is referred to by means of an expression that occurs outside the scope of the belief-operator. The satisfaction conditions of John's belief as reported in (19) are exactly the same as:

(13) John believes that Superman flies

Since Superman is identical with Clark Kent, John's belief is satisfied (correct) just in case Clark Kent (i.e. Superman) flies. But (13) captures what is genuinely intentional about John's state, in the sense that the intensionality-with-an-s of (13) is a reflection of John's subjective perspective in having the belief that he does. Thus, although extensional *reports* can be perfectly legitimate, the most fundamental way of construing the *contents* of one's intentional states should be in terms of Fregean modes of presentations, rather than in terms of satisfaction conditions extensionally expressed. Crane captures this well:

...there are ways of ascribing or attributing beliefs which relate the believer to the object of belief, creating an extensional context. But it is consistent with this to say that the belief which is ascribed is nonetheless a relation to a complete and Fregean proposition. The fundamental reason for maintaining this derives from the conviction that, in any state of mind, its intentional object is presented (thought about, desired, etc.) under some aspects, to the exclusion of others. So a full description of such a state must attempt to capture these aspects. (Crane 2001: 117)

To sum up: I have articulated a conception on which intentionality is direction on the world from a certain perspective, and hence a conception on which the problem of intentionality is the problem of understanding how we can be so directed, and what it means to say that we are. I've also argued that the best way to understand *direction* is by means of a representational conception of intentional states. In particular, the notion of content is explanatorily useful in two ways: first, it provides, together with the notion of schematic object, a way to understand how we can be directed on objects that do not exist – and thus cannot contribute to the individuation of a relation. Second, it enhances our understanding of what it means to have a perspective, and thus of how sameness of attitude and object still leaves room for epistemically and phenomenologically different ways of being directed on the object.

With this conception of intentionality in place, we are finally in a position to approach the problem of perception. In the next section I explain why I think that the representationalist picture of intentionality applies to perception, and then I proceed to examine how one might disagree with such an application, regardless of what picture of other intentional states one endorses. This will provide the foundations for the discussion of the intentionality of perception that I will conduct in later chapters.

4. Perception

4.1 Perception and Intentionality

This is a work on the intentionality of perception. More precisely, it makes a case for the claim that perception is intentional in a sense that is importantly similar to that in which paradigmatic intentional states, such as thought or belief, are intentional. It is an interesting

question to what extent a parallelism between perception and other intentional states can be maintained without overlooking significant peculiarities of perception, and it will be my concern to specify these as this work progresses, with particular attention to the relational character and phenomenological properties of perception. At first blush, there are some intuitive grounds for thinking that the intentionality of perception can be construed in the same way as the intentionality of other states, where this implies, on the conception that I have articulated, at least the following three claims: (i) every perceptual state has an object (ii) there are some perceptual states whose objects do not exist (iii) every perceptual state has a representational content. Elaborating on these points will give us a first approximate picture of our subject matter, and a way to set up the dialectic for rest of the thesis.

To start with, it seems obvious that if any subjective condition is ever directed on the world, then perception is. Just as it is not possible to think without one's thoughts being about something, it is not possible to perceive without one's perceptions being of objects, properties, states of affairs and events.

It also seems clear that sometimes things are not as they appear to be in perceptual experience: in illusion, it might seem as if a ball is red although in fact is blue; in hallucination, it might seem that there is a red ball in front of one although there is no ball at all. Following the model of intentionality explicated above, it seems *prima facie* reasonable to construe hallucinatory experiences as intentional perceptual states whose objects do not exist. Furthermore, given that even when we do perceive an object, it might be the case that it is in fact different from the way it appears to be – i.e. the experience might be illusory - it seems to follow that perceptual experience has conditions of satisfaction in the same sense as other intentional states do. Thus Searle:

In both the cases of belief and visual experience I might be wrong about what states of affairs actually exist in the world...In the case of the belief, even if I am in fact mistaken, I know what must be the case in order that I not be mistaken, and to say that is simply to say that the Intentional content of the belief determines its conditions of satisfaction; it determines under what conditions the belief is true or false. Now exactly analogously I want to say that in the case of the visual experience, even if I am having a hallucination, I know what must be the case in order that the experience not be a hallucination, and to say that is simply to say that the Intentional content of the visual experience determines its conditions of satisfaction; it determines what must be the case in order that the experience

not be a hallucination in exactly the same sense that the content of the belief determines its conditions of satisfaction. (Searle 1983: 39-40)

Finally, it is also obvious that if anything is perspectival, perception is: we necessarily – and literally – perceive things from a particular perspective, and it is in the nature of perspective that some aspects of the scene are disclosed to us from our vantage point, and others aren't. We are always perceptually directed on a scene under certain aspects and to the exclusion of others, and this perspectival element is of consequence for the phenomenological aspects of perception, since, for example, seeing an object with a circular shape from above is not the same as seeing it at a skewed angle. Hopp (2011) makes a similar point by saying that we must distinguish between the *what* and the *how* of a representation:

When we are asked to describe the way in which Degas depicted women, we might give two different sorts of answers. If we are interested in how he depicted them as being, we might say “as dancers.” But if we are interested in the way he depicted them as being dancers, we might say “impressionistically.” This is not a property of what is depicted, but a property of the depiction itself. A clear photograph might represent those same women as dancers, and not represent them *as being different*, yet represent them *differently*. In short, we must distinguish the *what* of a representation from its *how* (Hopp 2011: 22).

Although Hopp writes this in the context of a discussion of content in general, I think some of what he says can be particularly useful in understanding perspective in perception. What Hopp says is that we should mark a distinction between representing things *as being* a certain way and *the way* in which things are represented *as being that way*. I think that the significance of perspective in perception lies in this distinction. For example, a uniformly colored object might be represented as being of the same color all over, although in different ways from different perspectives, as the illumination varies. In the next chapter I take up this theme in far greater detail, and argue that only approaches which conceive of perception as being intentional in the sense that I've articulated – minimally, as having content – can make sense of the significance of perspective in perception, and of how this enables us to experience concrete objects as being mind-independent.

Something along these lines has been the orthodox conception of perceptual experience at least since the eighties and until very recent years, and as late as 2010 Tyler Burge wrote that

“it is *undisputed* that...beliefs, perceptions, and so on...are representational” (Burge 2010: 62). In fact, the whole characterization that I have given of perceptual experience – except perhaps its being directed on the world – has been challenged by authors who think that if perception really is to be the most fundamental way in which the world is disclosed to us, then it *cannot* have representational content. It is such a conception of the way perception relates us to our surroundings that will be my polemic target throughout the dissertation.

4.2 *Openness to the World*

There is disagreement among philosophers about virtually every aspect of perception, but there are two things which seem to be agreed on (almost) all sides: perception provides (i) a direct and (ii) basic access to reality. That perception is direct means that when we perceive nothing stands between the world and us, either in the form of mind-dependent sense data or in the way of inferential processes carried out by us: in perceptual experience we, as Soldati (2012) puts it, “enter into direct contact with items in the external world” (29).¹⁰ Perception is also basic, in that it is the most fundamental source of our conception of the mind-independent world we live in and a source of “*immediate* justification, a kind of warrant that does not depend on, for instance, any further inferentially acquired justification” (Soldati 2012: 29—30).

These features of perception are sometimes implicitly referred to by saying that perception is *openness to the world*. Thus *openness to the world* is a theoretical gloss on what is arguably our ordinary, pre-theoretical view of perception as the direct and most basic access we have to reality – what kind of theoretical gloss we should give on the pre-theoretical view is the main question that the present work tries to answer, and also the point where the agreement among philosophers comes to an end. The remainder of this chapter is devoted to explicating two possible philosophical conceptions of *openness to the world*, thus laying the foundations for a thorough investigation of the intentionality of perception which will be carried out in the subsequent chapters. First, I expound a conception of openness that makes use of representational content – one that is in tune with the content-model of intentionality

¹⁰ There might be inferential processes carried out by our visual system, but that does not make perceptual experience an inferential matter in the same sense in which the conclusion of an argument becomes evident for us only after we have gone through the premises connected to one another in an inferential structure.

that I have articulated. Second, I turn to purely relationalist conceptions, whose proponents reject the application of the content-model to perception.

There are at least two ways of explicating perception as openness to the world from a representationalist perspective. One way consists in drawing a disanalogy between perception and thought, which Husserl explains as a central case of a more general distinction between *signitive* and *intuitive* intentional states: very roughly, signitive states are those that intend an object in its absence, as when one stands outside a room and thinks that there is a table in there. By contrast, intuitive states are those that purport to make an object present, as when one enters a room and sees a table in there. Thus “A signitive intention merely points to its object, an intuitive intention gives it ‘presence,’ in the pregnant sense of the word” (Husserl 2001: 233). What the exact boundaries of the category of intuitive states are is not easy to say, but for present purposes it suffices to remark that perception is the paradigmatic case of an intuitive state, for in perception, unlike in thought, an object seems “to achieve full-bodied presence, to be there *in propria persona*” (Husserl 2001: 137).

Husserl’s elaboration on the signitive/intuitive distinction is to some degree problematic though, for Husserl thought that the distinction had to be explained by the fact that “signitive representation institutes a contingent, external relation between matter [content] and representative content [sensations], whereas intuitive representation institutes one that is essential, internal” (243). This raises the question what kind of internal relation there would be between intuitive representational contents and the objects to which these refer, and Husserl’s answer is that is one of similarity:

Only those contents can be intuitively representative of an object that resemble it or are like it. Phenomenologically put: we are not wholly free to interpret a content *as* this or *as* that (or in this or that interpretative sense)...since the content to be interpreted sets limits to us through a certain sphere of similarity and exact likeness (Husserl 2001: 244).

I think everyone will agree that spelling out the idea of a similarity-relation between representational contents and the objects of perception would be a daunting task, and hence if one could retain the spirit of the Husserlian distinction between signitive and intuitive states

while avoiding the trouble that comes with similarity, that would be desirable.¹¹ A clever way of doing this has been suggested by Elijah Chudnoff, who characterizes the class of experiences which purport to make their objects *bodily present* as having *presentational phenomenology*, of which he gives the following definition, clearly inspired by the Husserlian distinction between signitive and intuitive states:

What it is for an experience to have presentational phenomenology with respect to p is for it to both represent that p and make it seem as if you are aware of a truth-maker for p .
(Chudnoff 2012: 55)

One might see that John is walking by having John and his walk in view, or by looking at the trail of footprints he's leaving behind. In the former case, in addition to one's being in a state with a representational content to the effect that John is walking, one is also sensorily aware of the truthmaker for the relevant proposition, and so one is in a state with presentational phenomenology with respect to that proposition. In the latter case, one still bears some relation to the content that John is walking, but does not seem to be sensorily aware of the corresponding truthmaker; the latter state has presentational phenomenology with respect to the proposition that there are footprints, but not with respect to the proposition that John is walking.

Thus one way of characterizing the fundamentality of the way perception opens our surroundings to view is with reference to *perceptual presence*: perception purports to make its objects present in the most direct possible way. But *openness to the world* is not necessarily to be elucidated by means of a contrast between perception and thought. In fact, *openness to the world* has been first popularized by John McDowell in the analytic philosophy by means of a close parallel between what one can think and what one can see:

That things are thus and so is the content of the experience, and it can also be the content of a judgement....But *that things are thus and so* is also, if one is not misled, an aspect of the layout of the world: it is how things are. Thus the idea of conceptually structured operations of receptivity puts us in a position to speak of experience as openness to the

¹¹ It is worth pointing out that Husserl certainly had in mind a structural, rather than pictorial notion of similarity, and that structural similarity has recently been used to defend a form of indirect realism by Sollberger (2013). Interestingly, a naïve realist such as Logue also appeals to similarity between the phenomenal character of perception and its objects in order to explain the sense in which perception gives us an insight into the nature of what we perceive – although she admits that it is hard to make the idea “maximally clear” (Logue 2012: 229).

layout of reality. Experience enables the layout of reality itself to exert a rational influence on what a subject think. (McDowell 1994: 26)

According to McDowell, perception is openness to the world in the sense that in perceiving one takes in aspects of the *layout of reality*, which can exert a rational influence on what a subject thinks because there is no gap between the kind of things that one can think truly, that one can see, and that are the case (i.e., these are all aspects of the world).

So there are two complementary ways of articulating *openness to the world*: one with reference to the phenomenal features of perceptual experience, and another with reference to its objects. Both McDowell and Chudnoff make explicit reference to content. However, that perception has content is precisely the claim that is denied by Brewer and Travis, and precisely in the name of *openness to the world*. So there must be some alternative conception of what it is to be open to the world in perceiving.

Let us distinguish, following Travis's lead (2007, 2013), between:

(a) *Things being as they are*

and

(b) *That things are thus and so*

The difference can be illustrated by means of an example:

A story. Pia's Japanese maple is full of russet leaves. Believing that green is the colour of leaves, she paints them. Returning, she reports, 'That's better. The leaves are green now.' She speaks truth. A botanist friend then phones, seeking green leaves for a study of green-leaf chemistry. 'The leaves (on my tree) are green,' Pia says. 'You can have those.' But now Pia speaks falsehood. (Travis 2008: 111)

Calling certain leaves green can be a way of speaking the truth on one occasion and a way of speaking falsehood on another. The leaves, just as they are on both occasions, are a case of *things being as they are*: they are what they are, and they don't change from the first case to the second. *That the leaves are green* constitutes a case of *that things are thus and so* – a truth-evaluable content which is true on one occasion and false on the other. The leaves being as they are is neither something true nor something false, so if we ever are to say something

true or false, we need a way to get from *things being as they are* to *that things are thus and so*. We need something that decides which aspects of things matter, and how, to the truth or falsity of content. This is for Travis a question of the particular occasion for saying something. It is a particular context which makes it so that a certain feature of the leaves, like their appearance, matters to the truth of a content to the effect that the leaves are green, and it is a different context or occasion which makes it so that other features, such as the chemical composition of the leaves, matter to the falsity of that very same content.

Things being as they are is the layout of reality. But *that things are thus and so* is also the layout of reality, because insofar as it is true that things are thus and so, there must be something that is thus and so. It would then seem that there are two senses in which one can speak of the layout of reality, and therefore two senses in which perception is openness to the world.

In which of the two senses of “the layout of reality” is perception a condition whereby we are open to the world? *That things are thus and so* is that for which the matter of truth arises. But the matter of truth arises only when something decides which aspects of *things being as they are* matter to the formation of content and to its truth value. We can get from *things being as they are* to *that things are thus and so* only when some aspect of *things being as they are* is selected as that which matters in a determinate way to the fixation of content. Now, if perception is the most basic contact we have with the world, then, one might say, perception presents us with the layout of reality only in the sense of *things being as they are*, because the selection of particular constituents of reality and of a way for them to matter to the truth or falsity of a certain content would not be quite as basic as required by the fundamentality of perception. Indeed, I take this to be Brewer’s and Travis’s line.

Thus Travis writes:

It is no part of what perception is — of how it opens our surroundings to our view — that in perceiving one is to appreciate one set of facts as to what things look like, and ignore others. (Travis 2004: 73)

And thus Brewer, in a similar vein, states:

Perceiving is not a matter of being saddled with representational content...It is rather a matter of the conscious presentation of actual constituents of physical reality themselves,

particular such things, just as they are, which is what makes all contentful representation of that reality in thought even so much as possible. (Brewer 2006: 172)

[Representationalism] trades direct openness to the elements of physical reality themselves, for some intellectual act of classification or categorization. (Brewer 2006: 174)

The distinction that Brewer draws between the presentation of actual constituents of reality just as they are on the one hand, and contentful representation of reality in thought on the other, seems to parallel Travis's distinction between *things as they are* and *that things are thus and so*, and the claim that a representationalist view of perception trades openness to physical reality for some act of classification clearly re-proposes the thesis that perception can be the most fundamental contact we have with the world only provided that it be not a way of selecting which aspects of reality matter, and which don't, to the formation and truth of content.

Distinguishing between two senses of "the layout of reality" and two corresponding senses in which perception is openness to the world makes it intelligible how one may want to hold that perception is intentional in the minimal sense that we perceive mind-independent objects, and thus in the sense that in perceiving we are directed on the world, but not in the more demanding sense that it has content.

I will come back to the question of *openness to the world* and to the passages from Travis and Brewer quoted above over and over again throughout the thesis because I think that those passages contain the essence of the disagreement between representationalist and purely relationalist views on perception, and that they express a fundamentally misguided conception of the way in which perception discloses the world to our view. Whatever merits the distinction between *things being as they are* and *that things are thus and so* may have with respect to programs of research in semantics and pragmatics, I don't think that the conception of openness to the world that results from that is adequate.

In the next chapter, I argue that there are distinctions between different facts as to how objects look, related to the perspectival nature of perception, that are difficult to capture from a purely relationalist standpoint. In the third chapter, I argue that most purely relationalist

views have problems to account for illusion, and in the fourth chapter that they have problems in explaining the phenomenal and epistemic determinacy of perception.

Chapter Two: **The Way Things Look: a Defence of Content**

Abstract

According to *representationalism*, perceptual experience is endowed with representational content: it represents the world as being a certain way. In the first section, I discuss in detail an argument given by Charles Travis (2004) against perceptual content (1.1, 1.2), with particular attention to its connections with current debates on look-statements (1.3) and self-knowledge (1.4). In the second section, I examine the phenomenon of perceptual constancy (2.1) and I argue that it suggests an important phenomenological distinction that can be deployed for a defence of content (2.2). I also support my specific conception of content by means of an argument that links content to the way in which the mind-independent nature of concrete objects becomes manifest in perceptual experience (2.3). Finally, I address two possible objections (2.4).

Introduction

One of the liveliest current debates in the philosophy of mind concerns the nature of perceptual experience and more specifically the way in which perception is directed toward the world. According to a well-established view, perceptual experience is representational, that is, it represents the world as being a certain way; let us call this *representationalism*, or *intentionalism*, or also the *content view* (Burge (2010), Byrne (2001, 2009), Crane (2001, 2006, 2009), Dretske (1981, 1995), Searle (1983) and Siegel (2010)). Mike Martin offers a canonical characterization of the view:

On such a view, perceptual states represent to the subject how her environment and body are. The content of perceptual experiences is how the world is represented to be. Perceptual experiences are then counted as illusory or veridical depending on whether the content is correct and the world is as represented. (Martin 1994: 464)

Three main claims typically endorsed by representationalists figure in the passage: a) the metaphysical claim that perceptual states represent; b) the claim that they represent in virtue of having a certain content; c) the claim that content has an explanatory function to play with

respect to illusion and hallucination: it's the semantic properties of content (its correctness or incorrectness) that make a given perceptual experience an instance of veridical perception rather than an instance of illusion or hallucination.¹

The idea that perception is representational has enjoyed large consensus at least since the eighties, but has recently been challenged some philosophers. Perceptual experience, these philosophers say, is, at least when veridical, relational, *instead* of being representational: it is fundamentally a question of a subject being perceptually related to mind-independent objects; let us call this view *pure relationalism* (Brewer (2006, 2008, 2011, 2015), Campbell (2002a, 2002b), Fish (2009) and Travis (2004, 2007, 2008, 2013)). Thus Bill Brewer says:

Perceiving is not a matter of being saddled with representational content...It is rather a matter of the conscious presentation of actual constituents of physical reality themselves, particular such things, just as they are, which is what makes all contentful representation of that reality in thought even so much as possible. (Brewer 2006: 172)

And thus Susanna Schellenberg (2011) characterizes the view (without endorsing it) explicitly in terms of relations:

Perceptual experience is not representational, but rather fundamentally a matter of a subject being perceptually related to mind-independent objects, properties, events, or the event in which such relations obtain. (Schellenberg 2011: 714-715)

Both semantic and metaphysical claims figure in these passages: the semantic claim is that there is no content in perception ("perceiving is not a matter of being saddled with representational content"); the metaphysical claim is that perceptual states are fundamentally relational (perception is "fundamentally a matter of a subject being perceptually related to

¹ I am setting aside the problem of veridical hallucination for now: it seems possible to hallucinate a red ball while standing in front of a red ball; now, for all that has been said in the main text, such an experience would have a veridical content, and yet it would be a hallucination, so it would seem that the veridicality of content alone cannot suffice to distinguish hallucination from veridical perception (see also Siegel's (2010) distinction between strong and weak veridicality (35—38)).

Later on (chapter three, section two), I will introduce the notion of singular contents and argue that representationalists need to incorporate singular contents in their theory. Singular contents make reference to particular objects, but since in veridical hallucination the experience does not refer to any particular object, singular contents are sufficient to distinguish between veridical perception and veridical hallucination: in veridical perception, the content is true of the object that the experience refers to; in veridical hallucination, the content is true of *an* object that simply happens (i) to be in the same region of space as that at which the subject has the impression of perceiving something, and (ii) to have the same features as those which are represented in the subject's experience. However, in veridical hallucination the object is not referred to by the experience.

mind-independent objects”). The two claims do not necessarily go together, in that one might hold that perception is fundamentally relational *and* contentful as Logue (2014), McDowell (2013) and Schellenberg (2011) do – this mixed view is indeed the kind of conception that I strive to develop in this work. However, there is a particular way of construing the way in which perception puts us in relation to the world – of how it “opens our surroundings to our view” (Travis 2004: 73) – that does imply a rejection of content. The idea is, in short, that perception simply confronts us with what is in our environment, by affording us sensory awareness of it; whatever representational content follows from that awareness is the result of the way we represent things to ourselves on the basis of that sensory awareness. One such conception of openness is endorsed by Bill Brewer and Charles Travis, and that is my polemic target throughout the thesis.

In the first section, I explore in detail the argument that Travis (2004) offers against content in his influential “The Silence of the Senses”, and I discuss the argument’s relation to important topics such as the logico-linguistic structure of look-statements, externalism about content, and self-knowledge. In the second section, I show how the argument can be resisted, by drawing on a phenomenological analysis of perceptual constancy.

1. The Silence of the Senses

1.1 Setting the Stage

The portrait of representationalism that we have seen in the introduction to this chapter was painted with a broad brush: perceptual experience has content, and it is in virtue of the semantic properties of content that any one given perceptual experience counts as veridical or falsidical. I now want to discuss in more detail what is required by the idea of representational content, with particular emphasis on the relation between a perceiver and the content of her experience. There are four features that, according to Travis, are constitutive of a representational conception of perceptual experience as such:

- 1) *The representing in question is representing such-and-such as so* (Travis 2004: 58). This serves the purpose of distinguishing the relevant sense of “representation” from (i) a sense in which P represents Q by being a trace, or an effect of Q, as when “a

ring on a tree represents a year's growth" (58), and from (ii) a sense in which P represents Q by being a substitute for Q, as when a bit of plastic represents "an infantry division in a game of strategy" (58). The sense which is instead relevant to present purposes is one on which there is a way things are *according to an experience*, where things may actually not be that way. The notion of things being a certain way according to an experience is reminiscent of a family of metaphors by means of which perception is described as *telling* us, or *testifying* that things are a certain way (Austin 1962: 11) or *informing* us about goings-on in our environment the way a newspaper does (Siegel 2005), or *containing claims* about the world (McDowell 2009: 11), which contrast with other metaphors that characterize our senses as *dumb* (Austin 1962: 11), or *silent* (Travis 2004). If perception is representational, it *says* something, where what is said may or may not correspond to the truth. This sense of "representation" differs from both of the senses briefly mentioned above: it contrasts with the sense in which P represents Q by being an effect of Q, because that doesn't allow for the possibility of things not being the way they are represented; if representing is being a trace, then the representing cannot take place if the thing of which the representation is a trace doesn't exist. It also contrasts with the sense in which a bit of plastic represents an infantry division, because the latter is merely conventional: there isn't anything in a bit of plastic, taken by itself, which is such as to make the bit of plastic *say* that your opponent has acquired in infantry division.

2) *Perceptual experience has a face value* (59). When a person tells you something, her words have a certain face value at which they may or may not be taken: you may trust the person and treat what she said as true, or instead refuse to believe her; it's your decision. The same goes for perception, if it is representational: there's something in an experience which makes it so that, on the face of it, things are a certain way. But nothing forces you to trust your senses and form beliefs that treat the experience as veridical. If perception is representational, it must be possible for you to decide whether not to take what it tells you at face value.

3) *Being represented to is not autorepresentation* (60). Taking the bit of plastic you see to signify that your opponent has acquired another infantry division, or taking the footprints on the threshold to signify that someone has walked in, is to represent to

yourself that things are a certain way on the basis of your experience – this is what Travis calls “autorepresentation”. Autorepresentation is neither something that perception does, nor something whose face value can be accepted or rejected, for it precisely consists in taking things to be a certain way, and therefore accepting that they are so. If perception is representational, then the representing must be something “which represents things as a certain way, and which is so to be taken” (61) – this is what Travis calls “allorepresentation”.

4) *The relevant representing must be recognizable by us* (62). There must be some feature of an experience in virtue of which we are “able to appreciate what it is that is thus so according to it” (62). Roughly, if you ask me whether it is a brown table, or rather a green chair that my experience represents to me, or whether it is a straight stick rather than a bent stick, I should be able to give you a definite answer. The feature of the experience that enables me to do that might be, Travis conjectures, the way things look in having the experience: “the idea is that one could tell the representational content of an experience by the way, in it, things looked” (63). If perception is representational, then one has access to the content of one’s experiences in virtue of the way things look in having them; content is *look-indexed*.

Although in the end I will accept as constitutive of a representationalist position all of the four features identified by Travis, I think they all raise substantive issues, and it is certainly not evident, contrary to what Travis says, that they all are “non-controversially part of it” (58). In particular, in 1.4 I offer a detailed discussion of the fourth requirement on content – *the relevant representing must be recognizable by us* – which I think is especially controversial, and plays a fundamental role in Travis’s argument. For now, let us take for granted points one to four and let us see how Travis builds his argument.

1.2 A Dilemma for the Representationalist

If representationalism were true, then the way things look in having an experience would make recognizable for a subject the content of that experience. This has to be understood correctly though, for looks need not figure directly in the contents of perception. If, for

example, something I see looks like a lemon, its look serves the purpose of indexing a content to the effect that the thing *is* a lemon, not that it *looks* like a lemon; perception has an assertive character with respect to how things are, and indeed invites one to believe that things *are* a certain way, not just that they *look* that way. The *rationale* for thinking that perception is assertive with respect to the way things are can be explicated by means of two different kinds of considerations. The first concerns the explanatory role of perception with respect to belief: if perception is to give us *immediate justification* to form beliefs about our environment, then its content must be committal with respect to the way things *are*, rather than *look*. For from a content to the effect that something merely looks F, there can only be a mediated transition to the belief that the object is F – where the mediating role would be played by the assumption that if something looks F, then it's reasonable to think that it is an F. Whereas if perception represents that an object *is* a certain way, there can be a direct transition to the belief that the object is an F.

The second element has to do with the characterization of experience as having a face value. If a certain experience represents two lines as unequal in length, we have the faculty to decide whether or not to take it at face value. But if the experience just represents that two lines look unequal, it seems there's no possibility of distancing ourselves from the experience – for if the import of the experience is just that things look a certain way, then in things looking that way the import is implicitly accepted. Think, for example, of an assertion like “things are not as they look”. There is some cognitive achievement implicit in the assertion: it implies that there is something given in experience, where that, whatever it is, is rejected. Now, if what is given is that *a* is F, we can perfectly make sense of someone not willing to believe that things are that way. But suppose that what is given in experience is that *a* merely looks F. How can we reject this? In *a* looking that way, there's an implicit acceptance of what is given in the experience, otherwise it would make sense for us to say that things look F, although they don't look F – *that* is what would be required for us to distance ourselves from the experience, but that is contradictory.

There are two notions of looks that Travis discusses, neither of which he considers fit to index content. I devote the larger part of this section to the analysis of one notion, and I briefly comment on the other toward the end. On the first notion of looks, there is some object or feature of an object which bears resemblances to other objects or features of objects. Thus Fido might look like a dog, i.e., he might look like dogs look; Fido shares a *demonstrable*

look, to adopt Travis's terminology, with other dogs. Likewise, a shirt might look blue, i.e., the shirt might look like blue things look. Generally speaking, objects exhibit demonstrable, comparative public looks:

...something looks thus-and-so, or like such-and-such, where it looks the way such-and-such, or things which are (were) thus and so, does (would, might) look. On this notion, Pia may look (rather, very much, exactly) like (the spitting image of) her sister...The shirt looks blue (in this light) – as a blue shirt (so viewed) does, or might. The sun, at sunset, may look red. A van Meegren may look (uncannily) like a Vermeer. (Travis 2004: 70)

Suppose there is something in front of you which demonstrably looks like a yellow lemon. Then, by representationalist lights, that *lemony look* should be what determines content: your experience has the representational content that there's a lemon because what you see looks like a lemon.² However, there are many other things the object in question looks like: a tennis ball seen from a certain distance, a yellow soap, an orange seen in certain light conditions, and so on. In looking like a yellow lemon, it also looks like all these other things and many others. We are thus faced with a problem: why say that the experience represents a lemon, rather than an orange, or any other relevantly similar object? The point generalises:

...in looking like *Y*, *X* also shares a look with many things...If Pia looks like her sister, she also, on some understanding, looks the way she herself does, so might, or would, look. On some understanding or other, she looks (just) like any of indefinitely many different things. There is thus a substantial problem. Which facts as to Pia's looking (like) thus and so matter, and how, to how things should be to be the way they look *simpliciter*? Which looks, if any, matter to what is thus represented as so? And how? And why? (Travis 2004: 71-72).

Any one given object will have visual similarities to (will look like) a number of different things. To determine how the object is represented as being, a criterion will be needed to determine which similarities (which facts as to how the object looks) matter, and which don't.

² In this expository part, I take for granted that perception does represent high order properties like *being a lemon*, although many people claim it can only represent basic properties like *colour* and *shape*. Travis's argument is supposed to work on either construal of perceptual content, but in the second section of this chapter I will articulate my response by focussing on *shape* and *colour*. For more on the topic of the admissible contents of experience see Hawley & Macpherson (2011).

I think the challenge the representationalist faces could be construed as a dilemma. On the first horn, one might say that all of the demonstrable looks of the object contribute to the determination of content. If one takes this route, then an object is represented as being each and every one of the things it looks like. Upon seeing an object *o* that is in fact a yellow lemon, the content of your experience might be as follows: [(*o* is a lemon), (*o* is a tennis ball), (*o* is an orange)...]. The obvious problem is that on this construal our experiences would be totally incoherent, for in being one of the things it looks like, the object cannot be any of the others. But since the truth of the content above would require of the object that it be a lemon, and a tennis ball, and an orange, it follows that the experience cannot possibly be veridical. And given that one might construct lemon-like cases for pretty much anything, the conclusion would be that none of our experiences could ever possibly be veridical.

One may want to avoid incoherence by supposing that perceptual experience represents an object as being one of the things it looks like, *or* another, *or* yet another, and so on. Upon seeing a yellow lemon, your experience would have a disjunctive content to the effect that there's a lemon, *or* an orange, *or* a tennis ball, and so on, where the content would be veridical just in case one of the disjuncts corresponds to how the object is. This move strikes me as desperate, for at least two reasons. First, I can't see a principled way to put an end to the chain of disjuncts. If something looks like a yellow lemon, it also looks like a golf ball in a certain light, or an orange, and so on indefinitely. That this kind of content is problematic can be brought into relief by reflecting on how we conceive the epistemological function of perception. We want perceptual experience to play a definite role in rationalizing belief and action, and a standard representationalist approach holds that a certain experience rationalizes a certain belief (or action) instead of another in virtue of its content. Seeing a lemon rationalizes, if your experience represents to you that there's a lemon, your cutting it in half and squeezing its juice on the escallop. But if the experience represents that there's a lemon, or a golf ball, or an orange, that doesn't rationalize your cutting the thing in half any more than your taking your golf equipment and start practicing in the kitchen with the lemon.

Second, if you are a representationalist, you want the contents of experience to be phenomenally accurate, that is, to respect how it is for a subject to have a certain experience. But disjunctive contents don't satisfy this condition. Upon seeing a yellow lemon, it normally doesn't seem to a subject as if she's looking at a lemon, or an orange, or a tennis ball. It normally seems to her that she sees a lemon, full stop. A theory of perception should account

for a certain degree of determinacy in the phenomenology of experience, but the idea of disjunctive contents goes exactly in the opposite direction.

On the second horn, one might say that in perception a specific look, among those that the object has, is selected as the one which indexes content (a lemony look rather than an orangy look, say). However, it's not clear how that could be accomplished, given that both are looks that the object has, and are, as it were, on a par. What is the experiential criterion on the basis of which one could say that the experience represents the object as a lemon and not as an orange, given that the object in question looks like both? There seems to be none, and therefore there seems to be no principled answer to the question raised by Travis: "Which looks, if any, matter to what is thus represented as so? And how? And why?" (71-72).

One way to challenge the second horn would be to claim that there are some non-perceptual factors that at least partially determine perceptual content. The selection of a particular look would be grounded in factors such as background knowledge, assumptions, inclinations and sensitivity to the particular occasion on which we see something. If we open the fridge and see a yellow object, for example, we would be more likely to have an experience whose content is indexed by the lemony look rather than the tennis ball look. Although I don't see an obstacle in principle for this response, it certainly begs the question against the picture of perception as openness to the world as understood by the brand of relationalism under discussion. According to Travis, it is "no part of what perception is — of how it opens our surroundings to our view — that in perceiving one is to appreciate one set of facts as to what things look like, and ignore others (2004: 73)", and for Brewer selecting a look over another would be to trade "direct openness to the elements of physical reality themselves, for some intellectual act of classification or categorization (2006: 174)", for registering certain set of similarities rather than others is, on this view, part of our conceptual classificatory engagement with the world, rather than part of genuine perceptual presentation of reality. That perception is openness to the world means, on this reading, that it puts us in relation with the here-and-now, without informing us that things are thus and so.

For the moment, I prefer to "play by the rules" and conform to the purely relationalist conception of openness, thus conceding the point that selection of a certain look kicks in at a second, cognitive stage. However, in 2.2 I challenge the purely relationalist conception of openness, and in 2.4 I argue that at least some non-perceptual factors, such as expectations, play an important role in the determination of perceptual content.

If we follow the purely relationalist insight on openness, we see how the second horn of the dilemma leads to an arbitrary decision: the content of your experience is that there's a lemon just because you say so *ex post*, that is, after you have gone on to believe that there is a lemon, or have selected a particular look. The outcome of Travis's argument, then, is that the content of experience cannot be "read off of the way, in it, things looked" (69). On the assumption that perceptual content is determined by the way things look (on the present notion of looks), it follows that perceptual content should just be abandoned.

The second notion of looks that I mentioned at the start of this subsection is a matter of what can be picked up from certain visual features of a scene and inferentially taken to indicate that things are thus and so, like when one says that it looks as if Ronaldo is going to score (that is, judging from the speed of his run, the uncoordinated setup of the defenders, the inexplicable fact that the goalkeeper is off on vacation...):

The second notion...is very much a matter of what can be gathered from, or what is suggested by, facts at hand, or those visibly (audibly, etc.) on hand. So it cannot look as if X on this notion where it is perfectly plain that X is not so. Further it looks as if X only where one has not actually seen, or observed (for himself) that X is so; in which case there would be nothing to gather. (Travis 2004: 76)

The apparent advantage for the representationalist in embracing this notion of looks is that one avoids the problem of incoherent contents:

On this second notion, if it looks as if Pia's sister is approaching, or as if the painting is a Vermeer, then there is a way things should be, *simpliciter*, for things to be the way they thus look: Pia's sister should be approaching; the painting should actually be a Vermeer. (Travis 2004: 76)

With demonstrable looks we encountered the problem that they "point in no one direction" (Travis 2004: 72), in that it is not possible to identify *one* way things should be in order to be the way they look. Instead, we don't have this inconvenience with the second notion. If it looks as if Ronaldo is going to score, then there *is* a way things should be in order for them to be as they look: Ronaldo must score!

What makes this notion unfit for the job, however, is that something of the form *it looks as if Ronaldo is going to score* is not the sort of thing that can determine perceptual content,

but rather the sort of thing we represent to ourselves on the basis of what perception affords us, that is, it's a question of what we make of what perception brings into view. From a purely relationalist perspective, perception only brings Ronaldo, the defenders and the goalkeeper into view. If it looks to us as if he's going to score, it's because we take certain elements to indicate that he's going to score. Perception does not represent that he will score; rather, we represent to ourselves (on the basis of the elements that perception affords us) that he will score; so indexing perceptual content on this second notion of looks would violate the condition that perceptual content be allorepresentation, not autorepresentation.

Though for different reasons, it seems as though neither of the two notions of looks under examination can give the representationalist what they need: in the case of demonstrable, comparative looks we cannot fix a specific way the world should be in order for the experience to be accurate, while in the case of epistemic looks inferences and autorepresentation are in play. In the next subsection, I consider two arguments for content based on semantic considerations, and I argue that neither of them is satisfactory.

1.3 *Phenomenal Looks*

It is a highly controversial question whether “looks” is ambiguous between different senses or uses, and, if so, how many of these there would be. The debate was initiated by Chisholm (1957), with the distinction between three different notions of *appear words*, marked as the epistemic, the comparative and the noncomparative notion³. On an epistemic notion, one might say that “a ship “appears to be moving”, or that “it looks as though”, or “sounds as though” it were moving” (Chisholm 1957: 43), where this implies that she who utters the sentence “believes, or is inclined to believe, that *x* is so-and-so”, and that she has “*adequate evidence* for believing that *x* is so-and-so” (44). Instead, when appear words are used comparatively, the locution “*x* appears to *S* to be...” may be interpreted as “comparing *x* with those things that have the characteristic that *x* is said to appear to have” (45).

Although he never explicitly states so, I think it's fair to say that Travis's notion of demonstrable looks and the second notion of looks that he discusses correspond to Chisholm's comparative and epistemic notion respectively. But Chisholm also thought that there was a

³ It is often difficult to determine whether the authors participating in the debate intend to make semantic or pragmatic claims, or both. Therefore, I will follow Travis's deliberate vagueness and simply speak of “notions” of looks.

third, noncomparative notion of appear words, of which we find no mention in Travis's work. Indeed, Chisholm thought that a sentence like "The mountainside looks red" could be ambiguous between a comparative reading, on which the sentence entails "the mountainside looks the way red things look in daylight" (50), and a noncomparative reading on which the sentence comments on a certain appearance without relating or comparing it to other things. This was seconded by Frank Jackson (1977), who argued for a phenomenal notion of looks – fairly close to Chisholm's noncomparative notion – in addition to an epistemic and a comparative notion. In particular, Jackson emphasised that this supposed phenomenal notion concerns basic perceptible properties such as shape, colour and size:

The phenomenal use is characterized by being explicitly tied to terms for colour, shape, and/or distance: 'It looks blue to me', 'It looks triangular', 'The tree looks closer than the house', 'The top line looks longer than the bottom line', 'There looks to be a red square in the middle of the white wall', and so on. That is, instead of terms like 'cow', 'house', 'happy', we have, in the phenomenal use, terms like 'red', 'square', and 'longer than' (Jackson 1977: 77)

There has been a revival of the debate concerning looks in the contemporary literature, and sometimes a tendency to move from questions concerning look-statements to conclusions concerning the metaphysical or phenomenal nature of perception. It certainly is puzzling that Travis should not even mention a phenomenal notion of looks, and it might be tempting for the representationalist to say that even if the reasons given against indexing content either on comparative or epistemic looks were conclusive, it would still be legitimate to doubt whether one was forced to rely on either of those notions in order to establish a representational view of perception.

It seems clear that epistemic looks are not fit for indexing perceptual content, and it also seems clear that comparative looks are, at least on Travis's understanding, *public looks*, that is, objective, public features of objects or visual scenes. If we say that an object looks like a lemon, we are committed to there being visual similarities between the object and lemons; it is a fact as to how the object looks that it has similarities to lemons, and it is a fact about the object that we report when we say that it looks like a lemon. However, one might think that it is not the way things are and publicly look which should be used in determining perceptual

content, but rather a phenomenal and subjective notion of looks, such as *looking thus and so to a perceiver* – or so Siegel (2010) and Byrne (2009) think. Thus Siegel writes:

...the question itself seems flawed driven as it is by the idea that demonstrable looks might fix contents of experience. If the fact that a lemon demonstrably looks lemony does not entail that it looks lemony to S, why should we think that the lemon's demonstrable looks fix the facts about the contents of S's experience when she sees a lemon? At best, these facts are fixed by S's experience somehow picking up the demonstrable look of the lemon, when she sees it. But with the notion of *picking up on* a demonstrable look we've introduced another kind of looking altogether. Demonstrable looks are irrelevant to fixing the content of experience. Since they are public, they are part of the way the world is and, as such, do not automatically determine how the world appears, seems, or is presented to a perceiver. If any notion of looking is going to constrain the contents of experience, it must be looking some way to a perceiver. (Siegel 2010: 62)

There are two remarks which Siegel makes in this passage. The first is that if perceptual experience somehow picked up on a demonstrable look as the look which indexes content, then the fact that a lemon has countless similarities to other objects would no longer be problematic, for there would be one look which determines content, and others that don't. The second is that if it were the way the object looks *to a subject* which indexes perceptual representation, then again there being a multiplicity of demonstrable looks would be irrelevant. On this second option – which seems to be the one approved of by Siegel – there is some notion of looks on which look-statements neither report on visual similarities among objects, nor do they report on elements in a visual scene that might be taken to indicate that things are thus and so (the case of the epistemic notion of looks); rather, they report on the way things look *to a subject*. If the lemon looks to S to be a lemon, then the content is something like *there is a yellow lemon*, or *that is a yellow lemon*; if it looks to S to be an orange, then the content is *there is an orange*, and so on.

Alex Byrne (2009) raises a related point in favour of there being a notion of things looking thus and so *simpliciter*, rather than comparatively:

...if cows look F, and that sculpture looks like a cow, then that sculpture looks F. Cows look to have a distinctive shape – cow shaped, for want of a better term. Given contingent facts about the way cows look, to look like a cow (comparative) is to look, *inter alia*, cow-shaped. What is *that* use of “looks”? Not comparative, on pain of a regress.

Apparently not the epistemic one either. In a distorting mirror, something might look cow-shaped but not look *as if* it is cow-shaped. (Byrne 2009: 441)

This is an argument from regress: looking like a cow is to look cow-shaped, but looking cow-shaped cannot in turn be understood in comparative terms, on pain of a regress. Therefore, we should allow for a notion of looks on which “looking F” is to be understood in noncomparative, or phenomenal terms. This would imply, in turn, that there is a specific way that perceptual experience represents the world to be, and hence, *pace* Travis, a definite set of accuracy conditions.

In fact, I find the arguments given by Byrne and Siegel not entirely convincing, and, generally speaking, I think that one should defend the idea of perceptual content by appealing to the way in which perception discloses the world to our view in virtue of its phenomenological aspects, rather than by relying on linguistic considerations. In particular, Byrne’s argument from regress seems to me to conflate matters of metaphysics with matters of semantics. If a sentence says that X is like Y in some respect or other, then for that to be true there has to be a truthmaker. Consequently, one might think that, as a matter of metaphysics, this has to be some property F other than that of *being similar in some respect*. If, for example, “the ball looks red” means that the ball looks like things that are red, then, one might think, there has to be a property such as *redness* — or perhaps an appearance-property such as *redness** — which makes the sentence true. This much is a matter of metaphysics, though. It is a further, distinct claim to suppose that the property which verifies a comparative look-statement should itself be directly specified by another look-statement, in which “looks” occurs in a noncomparative or phenomenal sense. Byrne does not give any argument in favour of this additional constraint, and in absence of such an argument his considerations are incomplete.

Siegel’s argument is also inconclusive, for it is controversial whether *things looking some way to a perceiver* picks a notion of looks fundamentally distinct from Travis’s demonstrable looks. Mike Martin (2010) has articulated at length an account of looks on which a fundamental explanatory role is played by a comparative understanding of look-statements, and on which statements of the form “o looks F to S” are analyzed on the basis of the more basic form “o looks F”, with “S” figuring in the more complex form as an adjunct and modifier of the verb “looks”, rather than as making explicit an argument-place that would be

already implicitly present in the simpler “*o* looks F”. Martin labels his own account *Parsimony*, because of the methodological and metaphysical parsimony that characterizes it:

Parsimony: looks statements are made true just by properties of objects that we need to appeal to in order to explain the truth of sentences that are not explicitly looks sentences.
(Martin 2010: 197)

Parsimony says that the properties that we are committed to attributing to objects anyway in our discourse and through our practices are sufficient to explain the truth or falsity of look-statements, and hence that there is no need for positing special appearance-properties. In addition, Martin thinks that “it is common knowledge among us that others who have normal vision know the looks associated with properties they know of through vision” (195), and that we exploit this common knowledge in our ordinary talk of objects and their looks. When it comes to understanding look-statements, this common knowledge proves crucial, as she who understands a statement to the effect that *o* looks F can be imagined as comparing *o* and its look with the look of F-things, of which she has knowledge. She who understands a statement like “that model looks pregnant”, for example, can be imagined as executing “a function *getting-the-characteristics*” (169), which returns for the lexical item “pregnant” a suitable property, *pregnant*, with respect to its look and relative to a contextually determined metric of similarity *k*; what the sentence conveys, then, is that the model has a certain look, which is similar to the characteristic look of pregnant women.

Then Martin makes a further step, which consists in assuming that the comparative structure is “reflected in logical form” (172). He also suggests the following formalization, based on a Davidsonian⁴ approach to the structure of action sentences and eventives:

$$\exists s [\text{has}(\text{that model}, s) \wedge [\text{look}(s) \wedge \text{SIM}(\text{C}(\text{pregnant}, \text{look}, k), s)]] \quad (172)$$

This is to be understood as a quantification on a state (or event) of looking, which is “had” by the model. “SIM” (for similarity) is a predicate that takes as arguments a predicate F (pregnant) and an individual (the model), and returns the value TRUE just in case the individual has the characteristic look of the Fs. “C” is in turn to be understood as the function

⁴ Davidson (1967), Higginbotham (2002).

getting-the-characteristics, which returns “the property characteristic of the kind “pregnant”, with respect to “looks” when we are considering the understood restriction, human female in the n th week of pregnancy” (172-173).

There seems to be no reason why one should not extend this proposal to sentences involving terms for colours, shapes or sizes. Building on Martin’s analysis, the logical structure of a sentence like “the ball looks red” might be rendered as follows:

$$\exists s [\text{has}(\text{ball}, s) \wedge [\text{look}(s) \wedge \text{SIM}(\text{C}(\text{red}, \text{look}, k), s)]]$$

That is, the ball has a look which is similar to the characteristic look of red things, given a contextually determined metric of similarity. So, it seems the comparative analysis of looks can be extended to statements that report on basic perceptible properties such as color. At this point, one might remark that Siegel’s objection was that the way things look *to someone*, rather than just things having a demonstrable look, is relevant for content, and that Martin’s analysis is silent on this crucial point. However, if Martin’s account is viable in the first place, then there doesn’t seem to be anything preventing one from conceiving of statements of the form “o looks F to S” on the basis of the comparative analysis of the simpler “o looks F”, with the addition of an adjunct (“to S”) that introduces a partial modification in the meaning of the sentence, in the same way in which, in a sentence like “John lit a cigarette in the park”, the expression “in the park” partly modifies the meaning of the simpler “John lit a cigarette”. Although Martin does not offer an explicit formalization, we might volunteer the following logical analysis of “The ball looks F to S”:

$$\exists s [\text{has}(\text{ball}, s) \wedge [\text{look}(s) \wedge \text{SIM}(\text{C}(\text{red}, \text{look}, k), s)] \wedge \text{Experiencer}(s, S)]$$

That is, the state or event of looking in question is one that has the similarities described above, and in addition is experienced by S. It is worth noting that Travis himself very briefly mentions the subject of look-statements which include reference to a subject:

‘Looks like’ takes modifiers. Pia may look like her sister. She may also look like her sister, viewed from this angle, or in this light. Again, she may look like her sister to Sid (viewed from this angle, or in this light...If Sid is, say, bad at facial recognition, she may

look like her sister to Sid, though she does not really look like that at all. But all the more will there be countless other things she looks like to Sid — Zoë, no doubt, for one. Looking like X to N takes account of particular (I think contextually fixed) ways in which things may look different to N than they actually look — inebriation, fatigue, colour blindness, etc. But these ways do not change the fact that if X looks like one thing to N, then it looks like countless things — on the notion of looks being set out here. (Travis 2004: 70, fn. 12)

Although this is quite sketchy, it is enough to invite two considerations: first, Travis seems to embrace something in the spirit of Martin's account, not only in that he offers a comparative analysis of look-statements, but also in that he treats the specification of an experiencing subject as a modifier. Second, he shifts the focus from semantics to the notion of openness to the world, for he suggests that even if content were indexed on a subjective notion of looks, it would still be the case that there are several subjective looks, and that each of them has right to be *the* look which indexes content — unless, of course, one could provide a criterion for the principled selection of one look over the others.

This ties in with the principal points of *Parsimony*: although *Parsimony* by itself does not prove that there *cannot* be a distinct phenomenal notion of looks, it certainly does lessen the impression that there should be something special about look-statements concerning terms for shape, size or colour, or that the form “*o* looks F to S” should be understood as fundamentally different from the simpler “*o* looks F”. In addition, construing look-statements that include explicit reference to a subject along the lines of *Parsimony* offers an explanatory advantage, in that it casts light on the fact that the statement “The ball looks red to S” is both about a psychological state *and* about a ball, just as the simple “The ball looks red” is about a ball (and its look). That is, construing *looking F* and *looking F to S* as points on a continuum, rather than different kinds of events or notions of looks, allows one to see how the knowledge that we have of the characteristic look of things, and that we exploit in conveying information to others about those things (or their looks), is not lost or altered when one moves from “*o* looks F” to “*o* looks F to S”; the kind of information conveyed by the simpler form is preserved in, and also conveyed by, the complex one — this suggests, in turn, that the very same notion of looks must be in play.

Siegel seems to be in a predicament: if she is assuming that *looking F to S* picks out a fundamentally different notion than Travis's *looking like F* (which she certainly is), then

Parsimony seems to undercut her argument. If, on the other hand, she were to accept a comparative understanding of *looking F to S*, then this would fall back into the orbit of Travis's argument, as it were, for there would be several things that any one object would look like *to S*. All things considered, it seems that supporting representationalism either by means of logico-linguistic considerations or by setting up a supposed explanatory regress is problematic.

The representationalist might then change their tack and target the controversial requirement that perceptual content should be accessible to the subject of the experience. In the next paragraph, I examine several externalist strategies with respect to fixation of content, and argue that representationalists should endorse the requirement on accessibility after all.

1.4 *Content and Accessibility*

A lot in Travis's argument depends on the condition that the content of any one given experience be recognizable by, or reflectively accessible to, the experiencing subject *on the basis of the experience* itself (let us call this condition *accessibility-requirement*). However, one might think that mental content is externally determined: there are factors that decide what way things are represented as being, where these factors cannot be appreciated on the basis of the experience, for they have their origin in evolutionary facts about our sensory system and other relational factors involving our environment. Therefore, I must get clear on whether the accessibility-requirement is essential to the brand of representationalism that I want to articulate, and why. To this end, it will be fruitful to take a step back from Travis's argument and consider the question of the fixation and recognition of representational content in the larger context of the philosophy of mind of the last decades.

All theoretical enterprises based on the notion of representation face what is often called the *disjunction problem*:

This is the problem of explaining conditions on representation that show that representation applies to one set of attributes rather than equally well to a set of alternative attributes. Philosophers concerned with the problem ask why a perception that intuitively represents a body does not equally represent a light array. Or it is asked why a frog's perceptual representation that seems to be as of moving objects might equally well be as of flies, bee-bees, features, stuff, instances of abstract entities, undetached parts of

objects, or temporal slices of objects. The challenge is to explain conditions on representation that show why representations represent one range of entities rather than other entities that co-vary with, and in many cases play a role in causing, the representation. (Burge 2010: 322)

There is an indefinite number of things and circumstances, or so the idea goes, that prompt the frog to lash out its tongue: not just flies, but also black dots, undetached parts of flies, and so on. In the absence of a criterion for selecting one thing, from among the indefinitely many possible candidates, as *the representatum*, the project of providing a representational account of our mental life or that of other organisms seems bound to fail. In particular, it seems problematic to explain how representation can be *misrepresentation*. For suppose that a state represents As by being reliably caused by As, although sometimes is also caused by Bs. But then what it seems correct to say is that the relevant state represents *A or B* rather than simply *A*. Therefore, a B-caused token of the mental type in question would not constitute a case of misrepresentation.

There are several related, but distinct, strategies that one might take with respect to the disjunction problem, both for what concerns mental representation in general and for perceptual representation in particular. In what follows, I briefly expound two possibilities suggested by Fodor and Dretske respectively, and point out that, *prima facie*, there is a problem for them to satisfy the accessibility-requirement. I then examine Burge's account of self-knowledge and discuss whether his brand of externalism can be reconciled with a somewhat weakened version of the accessibility-requirement (and whether the resources of Burge's account can also come to the rescue of Fodor's and Dretske's respective theories). Finally, I discuss Travis's considerations in favour of the accessibility-requirement itself.

Jerry Fodor (to whom the label "disjunction problem" is due) has famously argued that the relevant causal relations mentioned above differ in their counterfactual properties, for there is an asymmetric dependence of the B-caused instances of our target mental state (or occurrence of a token symbol "A" of *Mentalese*, to assume Fodor's framework) on the fact that the token mental state is normally A-caused. Suppose that "horse", as a token symbol of *Mentalese*, is normally caused by horses, but sometimes also by cows. According to Fodor, the tokening of a cow-caused "horse" depends on there being a causal connection between horses and tokenings of "horse" in the first place, because misidentifying a cow as a horse wouldn't have led one "to say 'horse' *except that there was independently a semantic relation*

between 'horse' tokenings and horses” (Fodor 1987: 107). Further, the dependence in question is also asymmetric, because, “since 'horse' does mean horse, the fact that horses cause me to say 'horse' does not depend upon there being a semantic – or, indeed, any — connection between 'horse' tokenings and cows” (108).

The counterfactual solution offered by Fodor is of course ingenious, but doesn't put one in a position to address issues of content accessibility. In a particular perceptual case in which one sees a yellow lemon that looks both like a yellow lemon and a lemon-shaped soap, for example, one could only have access to a content to the effect that there was a yellow lemon in front of one, provided that one also had access to the truth of a counterfactual of the form *if there were no yellow lemons, my experience would not represent lemon shaped soaps as tokens of the same type as a yellow lemon*, and to the falsity of a counterfactual like *if there were no lemon-shaped soaps, my experience would not represent yellow lemons as tokens of the same type as lemon shaped soaps*. However, this faces two obvious problems: first, any access that one might have to the relevant counterfactuals could not be based on perceptual experience; second, perceptual content is supposed to be something the access to which can be easily gained by anyone, and thus not something that presupposes knowledge of a specific theory of content or counterfactuals related to it. So, whatever merit a Fodorian approach may have in solving the disjunction problem, it cannot (and it's not supposed to) be used for vindicating the accessibility-requirement.

The second attempt at solving the disjunction problem is due to Fred Dretske. Dretske (1995) emphasizes the importance of a teleological condition on representational properties of systems. The state of a speedometer might represent that a car is going, say, 30 kmh. Suppose the speedometer performs its function as follows: the pointer indicates 30 as a result of (i) the rate at which the axle is rotating and (ii) the height of the axle above the road (which provides, in turn, a measure of tire size). But since these are factors that co-vary with, and play a causal role in determining, the supposed representation of speed, why not say that the number 30 represents the rate at which the axle rotates, or the height of the tires relative to the road?

The answer suggested by Dretske is that systems represent a certain property in virtue of the function they have (or, in case of conventional representations, by the function they have been given):

The fundamental idea is that a system, S, represents a property, F, if and only if S has the function of indicating (providing information about), the F of a certain domain of objects (Dretske 1995: 2).

So, it is correct to say that the speedometer represents speed instead of the rotation of the axle or the height of the tire relative to the road because, although information about rotation and height is employed in, and co-varies with, the representation of speed, that information stops short of representational content, as it were; the speedometer has been designed to represent speed, not rotation.

What goes for speedometers goes for natural representations such as perceptual states, whose function is un-derived, or original, i.e., it has not been determined by agents already endowed with representational powers. In particular, the evolutionary story of, say, frogs and their sense-organs, together with the typical causes of their representational states, have made it so that their perceptual apparatus has acquired the function to represent bodies, and not, for example, arrays of light; information concerning arrays of light, although employed in the representation of bodies, stops short of content. What goes for frogs goes for us: during the period over which our perceptual system developed, there were no funny lights that made yellow look like orange, or additional hashes that made two lines of the same length look unequal, as in the Müller-Lyer diagram. Our sensory organs have evolved to detect the colour, shape and size of things in the conditions in which they were normally experienced, not under any logically possible condition. So it would seem there *is* a principled criterion to determine whether the content of a given experience concerns a lemon under normal light or an orange under yellow light: our sense-organs have evolved over our natural history in such way as to represent yellow in normal conditions rather than orange under unusual lighting.

At first blush, Dretske's teleological story seems to be affected by the same problem as Fodor's counterfactual account: those factors which determine the function and hence content of our perceptual states are normally not within our reach, and certainly not on the basis of experience. So, a teleological account solves the disjunction problem at the cost of dropping the accessibility-requirement. I should emphasize that this is an approximate assessment of Dretske's view though, and that I shall consider possible resources for the teleological theory toward the end of this section. For now, I simply ask the reader to bear with me.

Tyler Burge's approach combines a notion of representation other than that employed in informational models of content *à la* Fodor or Dretske with a particular version of the accessibility-requirement, on which content is recognised on the basis of factors other than the way things look. Burge labels informational models of representation "deflationary conceptions" (Burge 2010: 292), and he thinks that such conceptions, in the attempt of rendering the notion of representation scientifically acceptable, end up trivializing it, for they do not recognise the distinctive kind of psychological questions that representation helps to answer. According to Burge, philosophers who seek to reduce representation to information-theoretic concepts see themselves as saving representation from darkness "by reducing it to notions in sciences other than psychology, particularly natural sciences" (Burge 2010: 296), where in fact representation earns "its keep in science, and to a large extent in common sense, by figuring in successful explanation" (298).

Burge's insistence on the need to see representation as the essential ingredient of successful predictions concerning the behaviour of living beings is reflected in his discussion of the disjunction problem, which he thinks is "largely an artefact of reductive programs, detached from explanations in perceptual psychology" (Burge 2010: 322). Burge points out how several candidates for the role of *the representatum* in the case of a frog lashing out its tongue can in fact be easily discarded by taking into account the perceptual capacities of the organism, successful explanations of its behaviour, and by means of experimental setups. A perceptual capacity like constancy, for example, is explained only by supposing that the organism represents distal objects rather than proximal stimuli. More in detail:

What attributes are represented by a frog's vision? Philosophically contrived entities like undetached fly parts or temporal slices of flies are excluded from being represented as such by the fact that they are not kinds that ground biological explanations of the frog's needs and activities. I believe that undetached fly parts can probably further be ruled out by tests of frog attention. (Burge 2010: 322)

Even apart from artificial contrivances, it is often asked whether a frog's visual system specifies bee-bees, flies, light arrays, sense data, moving bodies, or what not. Most of these issues can be solved experimentally.

Given that the frog has a perceptual visual system, with capacities for perceptual constancies such as distance perception, there is empirical reason not to take the representations to be as of light arrays or sense data. (Burge 2010: 323)

According to Burge, difficulties in explaining away the disjunction problem and hence in accounting for misrepresentation derive from a misconstrued notion of representation, namely one that doesn't take seriously the scientific role of psychological explanations and doesn't pay enough attention to experimental results. Again, what holds for frogs holds for human beings: that we represent our environment as containing water is partly a matter of there being constitutive, non-representational relations between us and water (instead of *twater*, say). That we perceptually represent bodies, not light-arrays, can be appreciated by considering our capacity for perceptual constancy (more on this in section two), and that successful predictions of our behavior hinge on bodies rather than light-arrays. How does this brand of externalism fare with respect to the accessibility-requirement, though?

As we have seen, externalism in general faces problems concerning knowledge of one's mental states and their content, and Burge's version of it is no exception. Burge holds that in order to think that, say, there is water in our environment, there actually must be or have been water (not *twater*) in our environment. But once we add to this picture of mental content the intuitively plausible idea that we have first-person, introspective access to the contents of our own thoughts, it seems that we run into the consequence that we can come to know by introspection alone that there is water in our environment. However, this is surely implausible, for *that there is water* is an empirical proposition. One way to set up the problem is as follows:

P1: I can think that there is water in my environment only if there is (or was) water in my environment;

P2: I can know by reflection alone that I think that there is water in my environment;

C: I can know by reflection alone that there is (or was) water in my environment.

How does one reconcile externalism with the possibility of self-knowledge, given the obvious falsity of C?⁵ Burge thinks that we must distinguish between the *enabling conditions* of our first-order thoughts on the one hand, and the *entitlement* that we have to self-knowledge on the other. He explains this by means of an analogy with knowledge of logical and mathematical propositions. To achieve understanding and knowledge of a proposition like "nothing is both a dog and not a dog" (Burge 1996: 94), for example, one will need to have or

⁵ This is known as the McKinsey-problem for externalism; see McKinsey (1991).

have had some relevant experiences. But those experiences need not be part of the epistemic basis for believing the proposition in question; the warrant for believing that nothing is both a dog and not a dog is certainly not perceptual. Likewise for self-knowledge: there are enabling conditions for thinking that there is water in the environment; however, knowledge of those conditions need not be part of the entitlement that one has for knowing that one thinks that there is water. The first-order content to the effect that there is water is inherited by the second-order thought, without the need for the subject to be able to articulate the conditions in virtue of which the first-order thought has the content that it does.

One might try to apply Burge's account of self-knowledge to perceptual content: we represent bodies because there are constitutive, non-representational relations between states of our sense-organs and bodies, which determine the states in question as representation of bodies, instead of, say, light-arrays. Having access to the relevant relations is not a necessary condition on knowing that we perceive bodies, because the first-order content to the effect that there are bodies is inherited by the content of a thought to the effect that one sees that there are bodies. So, Burge's version of externalism would be compatible with a somewhat weakened version of the accessibility-requirement, on which perceptual content is not entirely look-indexed, but is still recognizable in virtue of enjoying relevant experiences and having first-person access to externally determined contents.

Something along the lines of Burge's account is applicable to Dretske's, and possibly Fodor's account. Indeed, Dretske makes a point which is similar to Burge's, only in terms of the distinction between knowing *what* one thinks or experiences on the hand, and knowing *that* one thinks and experiences on the other:

My first-person authority extends only to the facts that, *given that these external relations obtain*, are internally accessible to me — to content, to what I think and experience. (Dretske 2003a: 140)

Whatever way we have of telling what it is we think and experience is not a way of telling that we think and experience it. In the same way that if we know there is an external world, it must be in some way other than the ways we have of knowing what is going on in the world; if we know we have minds, it has to be in some way other than the way we have of telling what is going on in our minds. (Dretske 2003a: 140)

Given that there is water in my environment, that is, that certain external relations obtain, I may think on one occasion or another that there is water; this is *what* I think, and this is within the province of my first-person authority. If this is what I think, this entails that there is water in the environment, but my thinking a water-thought is absolutely not a way of coming to know that there is water in the environment; just as coming to know through perception what material objects the world contains is not, at least according to Dretske, a way of coming to know through perception that there is an external world. Likewise for access to content: one need not know facts about one's sense-organs evolution and function in order to know how things are according to one's experience. Rather, certain external relations determine content, and content determines the way things look to us, which is in turn something within the province of our first-person authority. Knowledge of the way things look to us is within the province of our first-person authority as much as the knowledge of what we believe, although the facts in virtue of which things look as they do (and the facts in virtue of which the content of our beliefs is what it is) are not.

There are important differences between Dretske's and Burge's respective research programs, especially in connection with the different notions of representation that Dretske and Burge rely on, but for present purposes it is better to emphasize two important commonalities, namely that (i) both accounts are subject to a McKinsey-style objection, and that (ii) both Burge and Dretske respond to that kind of objection by drawing a distinction between a subject's knowledge of what she thinks or experiences on the one hand, and a subject's knowledge of the conditions which enables her to have the thoughts and experiences that she does on the other.

Earlier in this subsection I said that on a teleological account content is not within our reach, but retrospectively we can see that that was ambiguous: if one means that the conditions which determine content are out of reach, then what I said is true; whereas if one means that the content and phenomenology which such external conditions determine is inaccessible, that is false. Given that only the latter kind of accessibility is required in order for externalism to imply a reasonable picture of self-knowledge, one might think that perceptual content can be externally determined and accessible to a subject after all.

However, I think that the externalist strategy under consideration is in fact quite odd. Consider the following example: "table" refers to tables, in virtue of certain causal connections and social practices. Had these connections and practices been appropriately

different, “table” would refer to chairs instead. Given that typically we do not have access to the relevant connections and practices, it seems legitimate to ask in virtue of what we are supposed to know that by means of “table” we talk about tables, not chairs. And the answer is ... in virtue of tables! This seems indeed the bare bone of Burge’s strategy: had there been twater in our surroundings, we would know our thoughts to be about twater. But since there is water, we know that we think about water; and all this without having access to the facts in virtue of which our thoughts really are about water – one does certainly not need to read *The Origins of Objectivity* in order to think about water.

The convoluted structure and the ineffectiveness of Burge’s picture has been exposed by Brewer in a passage where he comments on whether a subject knows that she refers to *redness* rather than to some other, structurally equivalent property:

Although this charade [Burge’s account] guarantees that the subject’s (pseudo-)self-ascription of a (pseudo-)belief about redness will be (pseudo-)true, it leaves him hopelessly ignorant about which truth this is supposed to be, which belief he has thereby self-ascribed. So the appearance that he thereby knows how he supposedly believes the world to be in the first place is just an illusion. Rather, the fatal ignorance at that first-order level is simply recycled at the level of the self-ascription itself. (Brewer 1999: 65)

The charade that Brewer mentions consists in the subject trying to tell herself at a second-order level what concept she uses, or what kind of thing she refers to, by re-deploying the word that she uses at the first-order level to denote the concept or refer to thing in question. Crucially, this cognitive routine is useless unless the subject already knows at the first-order level what concept she uses or what kind of thing she refers to. To use a different image, what the subject attempts to do sounds dangerously close to the case of the person who says “I know how tall I am” by putting her hand on her head: her “measuring” will *always* turn out to be accurate, even if her height were to change all the time, and even if she were unable to specify her height other than demonstratively; she knows that she is *that tall* just in virtue of the fact that she is *that tall*, just like we are supposed to know that we are thinking about water merely in virtue of the fact that there is water in our environment.

On balance, Burge’s attempt to articulate a sense in which content is accessible seems problematic. But what should we say about a view which dispensed with the accessibility-requirement altogether? Suppose that a type of experience was individuated in terms of its

typical causes, and that its content was determined by such causes: an episode of seeing a straight stick in water, for example, would have the content that the stick is bent, because that kind of experience is typically caused by bent sticks, even if this causal correlation is not something that experiencing subjects can appreciate on the basis of the experience alone. Would this kind of view be wrong, and why?

While sympathetic with the accessibility-requirement, I find the way in which Travis argues for it quite unsatisfactory, in that I think his reasoning is partly flawed and partly confusing. Travis's acceptance of the requirement is initially motivated by a brief remark which he makes commenting on a sense of "representation" on which P represents Q by being a trace or effect of Q:

If certain neural states, say, represent certain distal stimuli in being their effects or traces, or those yielding our awareness of them, that would not be to the present point. It would not amount to *their* representing anything as so; as if that were something they might do without its being so. For if for them to represent involves their being traces, then where there is no such thing for them to be traces of, they simply do not represent that. (Travis 2004: 58-59)

Travis seems to disallow the possibility of causal factors determining a relevant notion of content on the grounds that this would violate the principle according to which perception might represent things as so without things actually being so – for if perception represented x by being its effect, then it could not do so in the absence of x . However, this reasoning is clearly flawed. Let us suppose the identity theory of mental states is true, and let us identify a certain *type* of perceptual experience E with a certain type of neural event NE. Suppose, also, that NE represents x by being an effect of x . In order for this to be true, it is only required that there have been systematic causal connections between NE and x over a certain period of time – the period over which NE acquired the function of representing x . It does not require that, for any one given token of NE, there be x as a distal stimulus. Tokens of NE can be caused by instances of y or z , and still represent x , because x is the typical cause of NE. If Travis's argument for the accessibility-requirement is grounded in the idea that content being determined by external, causal factors would be incompatible with an experience representing things as so *independently* of whether things actually are so, then the argument should be rejected.

Other remarks that Travis makes are less than clear: he labels the idea of there being a way things are according to an experience *committed representation*, and is keen to stress that “only when there is committed representation can one be represented to” (62). But if this is to have any substance, an argument should be given for the following conditional: if certain neural states were to represent external configurations in virtue of a systematic connection between the two, that would not amount to committed representation. However, he doesn’t provide such an argument, but rather merely assumes the truth of the conditional when writing that “for such things [“representations” as used in vision sciences] to serve their explanatory ends, there is no need to see such representations as committed to anything being so” (59, fn.3)). Increasingly, the whole strategy becomes even more confusing when Travis writes that it makes “no difference to present arguments whether the representation that occurs in perception is committed or not” (62).

Despite the obscurity of Travis’s line of thought, I think there are other, more promising ways of defending the accessibility-requirement. One thing that should motivate the claim that one cannot “represent things to people as so in a way they simply cannot recognize as doing that”(Travis 2004: 63) has to do with the idea that if perception disclosed our surroundings to our view by representing things as so, it would have to bear its content on its sleeves, as it were; episodes of perception would have to represent *to us*, rather than just to our visual system. Perception certainly plays a role in our theoretical and practical deliberation, but how can it do so, if the way things are according to it are not accessible to us in the first place? In effect, Travis does *gesture* at something like this when he writes that “it should not come as a complete surprise someday, to be sprung on us by future neurophysiologists, that we are thus represented to (uselessly, of course, since we were all ignorant of it)” (86). If the conscious character of perception is not to be a mere ornament, then it should play a role in explaining the intentionality of perception and its relation to other attitudes. However, if the conscious character of perception is understood in terms of representational content (as it is normally the case in representationalist approaches), and content is not accessible by the subject, then the whole conception starts unravelling.

A second element that I think should be mentioned in support of the accessibility-requirement has to do with the face value of perceptual experience. As McDowell says, it must be possible to decide “whether or not to judge that things are as one’s experience represents them to be” (McDowell 1994: 11). For that to be possible, however, it must be

reflectively accessible, recognizable to one what it is that is represented as so. Suppose one sees the Müller-Lyer diagram: in order for one to have the possibility of distancing oneself from the experience, so as to be able to decide whether or not to treat it as veridical, it must be clear to one what it is that the experience invites one to take as true. So the accessibility-requirement is at least partly based on the point that experience has a face value. Experience having a face value, in turn, seems necessary to explain the phenomenon of disbelief in perception: upon seeing the Müller-Lyer diagram, one will most likely refrain from believing that the two lines are different on length. At least part one thought that might motivate the accessibility-requirement, then, can be articulated as follows: we are able to distance ourselves from our experiences by not taking them at face value, and this requires that what our experiences tell us be available to us. If we add to this the idea that content is look-indexed, we have it that the way things look in experience makes representational content available to us, where this, in turn, puts us in a position to decide whether or not to take the experience at face value. With *these* considerations adduced in favour of the accessibility-requirement, I would certainly be inclined to accept it as a constitutive feature of representationalism.

So far we have seen two dubious ways of responding to Travis's argument: the introduction of a phenomenal notion of looks and the rejection of the accessibility-requirement. In the next section, I articulate my own response to Travis, which is based on perceptual constancy. In particular, I argue that the phenomenon of constancy suggests a phenomenological distinction between two different ways in which facts as to how an object looks can contribute to determining content.⁶ This suggest, in turn, a revision of the purely relationalist conception of openness of the world.

⁶ That constancy might be a key factor in this debate was first suggested to me by Gianfranco Soldati during one of our discussions. Although I'm not sure that he would agree with the specific way in which I explain constancy, his suggestion has become, since that discussion, one of the main guiding lines for the present work, and I wish to thank him for that.

2. The Way Things Look

2.1 *Perceptual Content and Perceptual Constancies*

Perceptual constancy is that phenomenon in virtue of which we are able to see the intrinsic perceptible properties of objects despite variations in light conditions, distance and orientation relative to us. It is said, for example, that we can perceive “the uniform color of a yellow wall although parts of it are illuminated more brightly than others” (Schellenberg 2008:55), or that, walking toward a friend, it doesn’t look as though *she* is getting bigger; she looks to remain constant in size. And, the idea goes, “this is true despite its being the case that the size of the image that her body casts onto my retina does, of course, grow as I get closer to her” (Kelly 2010: 146). Something analogous is held to be true of shape: it is usually claimed that, in an ordinary sense, “a penny looks round both when viewed head on and when viewed from an acute angle, even though the area projected by the penny onto our retinae under these two conditions is very different” (Cohen 2014: 812).

Despite the intuitive, *prima facie* plausibility of these characterizations, it is notoriously difficult to give a phenomenologically accurate description of constancy. In particular, it is not obvious how to square two seemingly conflicting data:

- (1) We see objects as unchanging despite variations in viewing conditions;
- (2) Objects appear differently as the viewing conditions change.

In this subsection, I take as paradigmatic the case of shape, and use it to explicate in more detail the apparent conflict between the two data.⁷ In the next, I draw a phenomenological distinction that helps to dissolve the impression of a conflict, and also suggests a principled criterion for the determination of a recognizable perceptual content, thus providing a response to Travis’s challenge.

⁷ (1) and (2) also inform scientific research on constancy. In particular, the problem of constancy is often construed as a question of producing a computational mode for explaining how the visual system tracks intrinsic properties of objects while discarding transitory, relational features of a perceived scene. Matthen (2010) has recently suggested a different approach: Matthen thinks that also transitory properties like orientation and illumination are assigned a value by the visual system. For an up-to-date overview of the main scientific models, see Cohen (2014).

Suppose you look at a glass with a circular rim from an oblique angle. How does the rim look with respect to its shape? There are three possible options that I would like to consider: a) circular; b) elliptical; c) both. If one emphasizes perceptual constancy as “stability in perceptual responses across a range of varying conditions” (Cohen 2014: 814), then one will be inclined to say it looks circular, where this would amount, by representationalist lights, to the claim that an experience of the circular rim of a glass seen at skewed angle has a content to the effect that the rim is circular (on the assumption that content is look-indexed, that is).

A problem with this option is the following: there is a difference between the appearance of the rim when you see it from above (in such a way that the image it projects onto your retina is perfectly circular, say) and its appearance when you see it at a skewed angle; that much should be uncontroversial. But if the glass looks differently, and content is look-indexed, then we have to explain how it could look differently and yet be represented as having the same property on both occasions. This brings out the complex nature of perceptual constancy: it cannot *just* be a question of stable perceptual responses across varying perceptual conditions, for the elementary observation that there is a change in the way the rim looks contradicts this hypothesis.⁸

Saying that the glass looks, when seen obliquely, *circular from here*, and that that is the way the experience represents it, won’t help much, for in looking *circular from here*, it also looks *elliptical from here*; these are two facts as to how the glass looks, *from here*. This can be appreciated by imagining a borderline case, in which the glass with the circular rim is replaced with a glass built in such a way that, despite having an elliptical rim, it looks exactly like the other glass does (with respect to its rim) from the given viewpoint – this might happen if the elliptical rim was taller on one side, for example. But then, since the elliptical rim looks the same as the circular one, one might wonder why not say that the glass with the circular rim was represented as elliptical in the first place.⁹ To put it briefly: the “circular-option” does justice to (1), but not to (2).

Faced with these worries, you might observe the glass a bit more and decide it looks elliptical; after all, someone willing to draw a realistic painting of the glass from that angle

⁸ For discussion of this point, see Cohen (814 — 818).

⁹ This is more straightforward in the case of colour. One might say that the shaded parts of a table look the same colour as the parts that are more brightly illuminated. However, if the shaded parts actually were painted darker (to the right degree), they would look exactly as they do.

would have to draw it with an elliptical shape. The problem for this route is the converse of the one raised for the circular-option, namely: how should the rim look in order to look circular? It should look exactly as it does. To see that it is so might require a bit of reflection: one might be tempted to say that in order for the rim to look circular, it would have to look the way it looks when you see it from above. But this would fail to take into account that, because of the perspectival nature of perception, that is not at all how the glass should look from an oblique angle. Shifting the focus for a moment from shape to size might help to grasp the point: suppose you see someone from a moderate distance, and they look to be a certain height (180 cm., say). Next, suppose you see them a second time, when they have gone 15 meters farther. How do they have to look in order to look 180 cm? Certainly not the same as they looked on the first occasion, otherwise they would look significantly taller than 180 cm. Likewise for shape: part of what it is to see a glass as circular from the angle in question is to be presented with certain visual similarities that its rim has to things with an elliptical shape. The “elliptical-option”, then, does justice to (2), but not to (1).

If neither “circular” nor “elliptical” work as an answer to our question, “both” must be right. However, one might think this speaks in favour of the thesis of the silence of the senses: both the circular and the elliptical look are facts as to how the glass looks, and are as such on a par. Whether we take things to be one way or another cannot depend on representational, look-indexed facts, because it’s not by selecting one fact over another that perception discloses our surroundings to view. Depending on the angle of view, one might add, it would be more natural to pick up on one look or another (if you see the glass from above, for example, it would be more natural to compare it to a circular glass rather to an elliptical one); still, both its looking circular and its looking elliptical are facts as to how the glass looks, and are as such on a par.

According to Travis, the particular occasion on which we see something plays a crucial role in determining what would be appropriate to say about the look of things (on that occasion). However, he says, there is no route from occasion to content:

Which ways things looked on an occasion ... is, in general an occasion sensitive matter: there are questions whose (true) answers vary with the occasion for posing them ... On one occasion for describing a given perceptual experience, there may be definite things one *would* then say in saying X to look, or have looked, like Y – so, again, things would not. But what one would say on some one such occasion does not decide, of that

experience as such, how things looking as they then did relates to how things should be to be the way they then looked. (2004: 71)

In seeing the glass at a markedly oblique angle, one would be more inclined to say that it looks elliptical. But that doesn't decide whether the glass should *be* elliptical to look as it then does, and therefore does not determine the content of the experience.¹⁰ If Travis were right, the representationalist would face again the same dilemma discussed above: either say that the experience represents the rim at the same time as circular and as elliptical, or make a totally arbitrary decision as to which look is to determine content.¹¹

In the next subsection, I argue that that is not so, and that there is a principled way to determine a unique, reflectively accessible content.

2.2 Perceptual Content, Adverbial Modifications and Dynamic Perception

In this subsection and the next, I offer two closely related arguments which are mutually supporting: the first is an argument for the claim that different facts as to how any one object looks can contribute in different ways to the constitution of a coherent, unique and reflectively accessible perceptual content. The second is an argument for the claim that the notion of content helps us better to understand how the mind-independent nature of the objects of perception is made manifest in experience. In the process of articulating these arguments, I also go back to the theme of perspective that I developed in metaphorical terms in chapter one (section two) and I argue that the theoretical significance of perspective in perception lies in the fact that different perspectives on an object give rise to different modes of presentation of it.

The result of the analysis that I have conducted in the last subsection is that, when a circular glass is seen at an angle, the statements "the glass looks circular" and "the glass looks elliptical" are both true; given Travis's conception of openness, this would be a good reason to think that a corresponding perceptual experience of the glass does not have a determinate

¹⁰ For more on occasion sensitivity see Travis (2005).

¹¹ Restricting the range of represented properties to *shape*, *size* and *colour* already reduces, of course, the magnitude of the problem, for one doesn't have to explain why the thing is represented as a glass rather than, say, an unusual cover for an *abat-jour*. Since my reply to Travis is based on constancy, I will focus on basic properties to which constancy applies.

perceptual content. However, I think there is an important phenomenological distinction that Travis overlooks, or that his assumptions about openness prevent him from making. I would like to formulate the distinction as follows: it is the distinction between *what* shape (colour, size) is represented on the one hand, and *the way* in which that shape (colour, size) is represented on the other. Since we've been assuming that perceptual representation is a question of the way things look, we can reformulate both sides of the distinction explicitly in terms of looks. We thus get: the shape (colour, size) that an object looks to have on the one hand, and the way in which that shape (colour, size) looks on the other. When you see the glass at a skewed angle, that is, the rim does look *to be* circular, but its circularity appears in a different way than it does when you see the glass from above. Retrospectively, we can interpret a passage from Hopp quoted in chapter one (subsection 4.1) as suggesting a similar distinction. Hopp says that we should distinguish between representing things as *being* a certain way, and *the way* in which things are represented as being that way (Hopp 2011: 22). Thus a glass might be represented as being circular, although the way in which it is represented as circular involves an elliptical look.

The relation between different looks is phenomenologically more complex than Travis thinks. In particular, the distinction that I have drawn might be taken to suggest that there are two layers to perceptual content: on one layer, the circular look of the glass determines (indexes) the property that the glass looks to have. On the other layer, the elliptical look of the glass determines the way in which the circularity of the glass is made manifest to one from a particular angle – the elliptical look determines the mode of presentation of the circular shape of the glass. Further, the glass displaying an apparent elliptical shape is constitutive of what it is to have an experience of a glass as circular from the perspective in question.

More needs to be said about this supposed two-layered constitution of content, and on the role and nature of apparent shapes. Sometimes, locutions involving *appearance properties* are seriously meant to indicate that objects have, over and above their intrinsic properties, also apparent properties. Thus, Shoemaker (1994) adheres to the “Ways = Properties Principle”, and infers from the fact that an object can have, say, an elliptical appearance, to the fact that there is a real, nameless property that corresponds to that appearance. In a similar vein, Schellenberg (2008) talks of “situation dependent properties”, defined as a function of the intrinsic properties of an object and its situational features. By contrast, I employ the locution “apparent shape” (colour, size) simply as a convenient way of talking; on my view, changes in

appearances that fall within the limits of perceptual constancies are changes in *the way* something is represented, rather than in *what* is represented. In effect, as I explain below, I conceive apparent properties as mere adverbial, experiential modifications.

Let us call the look that indexes the represented property *ground look*, and let us define the other look as an *adverbial modification* of the ground look. This has to be understood as a functional distinction: the ground look is the look that comes to acquire the function of indexing the represented property, while the other look is the one that comes to have the function of modifying the ground look. I want to say, then, that when we are presented with both the circular and the elliptical look of the circular rim of a glass, in a situation that is within the boundaries of size constancy, the rim *elliptically looks circular*.

This particular way of understanding the “both-option” achieves two important things: first, it does justice to both (1) and (2), for it accounts for the fact that we keep on seeing the glass as being circular and for the fact that the appearance of the glass varies. By the same means, it also provides an answer to Travis’s challenge, for it allows the representationalist to avoid the first horn of the dilemma: there being several facts as to how an object looks does not imply that the content of experience is bound to be incoherent, for different looks can play different roles in the determination of content (to repeat: one look might determine the property that an object is represented to have, and another look the way in which that property is represented).

I acknowledge that much more needs to be said in order to make my view of content maximally clear, and I intend to do this by considering further challenges: first, I consider the objection that my account falls foul of the second horn, and I respond by articulating the concept of dynamic experience. Second, I explain how my account satisfies the accessibility-requirement on content. Third, I give further support to my conception of content by means of an argument that involves the mind-independent nature of the objects of perception and the way in which this is made manifest to us when we look at objects from different perspectives. Finally, I come back to the first horn of the dilemma.

One may object that my proposal falls foul of the second horn of the dilemma, for if the looks are on a par, there can be no principled criterion to decide which look indexes the property that an object is represented to have (ground look), and which one indexes the way in which the property in question is represented (adverbial modification). The two looks (circular and elliptical) might play either function interchangeably. Therefore, fixing the

content in one way or another is a matter of arbitrary choice¹²; talking of a function that a look comes to have implies that there is something which selects the look in question as having that function. However, this was exactly the pitfall of the second horn.

I think that this challenge can be met. Perceptual content is to be determined and accessed in the context of a *dynamic experience*, where by “dynamic experience” I mean an experience which comprises different views on the object. Suppose one veridically sees a glass with a circular rim. As the experience unfolds over time and from different perspectives, the glass will look to be circular, although the apparent shapes of the glass will be, from most points of view, elliptical to varying degrees, approximating to, or coinciding with, a circular shape from a limited number of vantage points. The experience of the circular glass, then, will be phenomenologically characterized by a specific dialectic between the shape that the glass looks to have on the one hand, and the way that shape looks on the other. If one were to veridically perceive an *elliptical* glass instead, there would be a *different* dialectic, because the ellipticality of an object is not veridically made manifest in virtue of an apparent circular shape.

Phrasing the point in terms of ground looks and adverbial modifications might prove helpful in clarifying this. We have seen that in a dynamic experience of a glass with a circular rim, one is presented with a series of elliptical apparent shapes approximating to, or coinciding with, a circular shape, in such a way that the rim, over the course of the experience, *elliptically looks circular*. The circular look being the ground look *explains* the successive adverbial modifications that we are presented with. Having the elliptical look as the ground look instead, would not explain the approximations of the elliptical shapes to a circular shape when the object is seen from above. Therefore, the prioritization of one look as the ground look is not arbitrary; it just is the correct explanation of appearances. Thus, the second horn can be avoided as well.

The sense of “explanation” that I have in mind is directly relevant to the satisfaction of the accessibility-requirement on content. On a longstanding tradition in the analytic philosophy of perception, perceptual experience is to be conceived of as the joint upshot of two factors: (i) mind-independent objects, persisting relatively unchanged over time, and (ii) our spatiotemporal routes through the world that are such as to afford us encounters with the

¹² The case of colour is similar. Consider a brown table with shaded parts. How do we determine whether it is represented as being uniformly brown with darker appearances in certain parts, or as being brown in certain parts and of a darker colour in other parts? As in the case of shape, there doesn't seem to be a principled answer.

objects in question.¹³ If I walk around the glass, for example, there is both a sense in which its appearance remains the same and one in which it changes, corresponding to its circular and elliptical look respectively. In order for me to see the glass as circular, it must be part of my stream of consciousness that the experience include certain adverbial modifications, and not others (that when I get to see the object from above, for example, there be circular or quasi-circular modifications). A determinate interplay between the ground look and its adverbial modifications, together with the persistence of objects as unchanging, play a fundamental role in helping us to make sense in the first person, *as perceiving subjects*, of the way our dynamic experiences of mind-independent objects unfold over time; indeed, we explicitly register this explanatory role when giving explanations concerning the way things appear, of the form “*o* appears *thus* because I see it *from this angle*” (*under this light, from this distance*).

That the sense in which the ground look explains the adverbial modification is an experiential, phenomenologically salient fact, can better be appreciated by considering the way in which the mind-independent nature of material objects is made manifest in perception when we see objects from different perspectives. In the next subsection, I give an argument for this claim, and relate it back to the question of content and the silence of the senses.

2.3 *A Sense of the World: Content and Mind-Independent Objects*

The point of the following argument is to give further support to the conception of content as two-layered, by showing that that conception best explains how it is that we get a sense of being aware of an objective, mind-independent world in perceptual experience, where the notion of mind-independence involves that neither the existence nor (at least a subset of) the (perceptible) properties of an object depend on a perceiving subject and her deeds, except where these interfere with the object and its features.

I now state the argument, and then turn to a defence of its premises:

¹³ The first explicit articulation of these points in the analytic tradition might be attributed to Peter Strawson, with his emphasis on the fact that the very notion of an objective world is related to determinate experiential structures, in that we could not “employ any ordinary empirical concepts of objects unless our manifold perceptual experiences possessed the kind of coherence and interconnection which is required for the application of such concepts” (Strawson 1996: 89). Correlatively, the structure of the experience is explained in terms of mind-independent, relatively unchanging objects, in that “we appeal to ordinary objects in understanding the actual and counterfactual course and nature of our perceptual experience of them from different points of view over time” (Brewer, unpublished: 1). For analogous lines of reasoning see also Brewer (2011, chapter 7), Campbell (2002b), Eilan (1997), Evans (1982, chapter 7).

P1: The mind-independent nature of objects is phenomenologically evident in perception;

P2: So there must be a feature of perceptual experience that is revelatory of this mind-independent nature;

P3: This feature is a duality of perception: things appear to retain their properties as the conditions of perception change, while at the same time appearing differently as the conditions so change;

P4. What explains this duality is that a particular look is selected as the look which determines what property an object appears to have;

C: Therefore, perception selects one from among different facts as to how any one object looks.

P1 says that material objects are perceived as being mind-independent. This is a claim that has appeal for a variety of approaches to perception, and certainly both representationalists and pure relationalists would subscribe to it. Indeed, it is very natural to explain at least some of the phenomenal features of our perceptual experiences in terms of stable, mind-independent features of perceptual objects. A pure relationalist author like Brewer emphasizes this:

We appeal to ordinary objects in understanding the actual and counterfactual course and nature of our perceptual experience of them from different points of view over time.

(Brewer, unpublished: 1)

In addition, the very distinction between after-images and perception proper is often cashed out in terms of the sense of a mind-independent world that we get in perception as opposed to the case of after-images. This sense of mind-independence is in turn explained, at least partly, in terms of the possibility of changing our perspective on perceptual objects, which one typically does not have with after-images and alleged cases of pure sensations. Thus Siegel emphasizes this point with respect to the spatial location of objects: “If S changes her perspective on *o*, then *o* will not thereby move” (Siegel 2010: 178). Along the same lines, Textor distinguishes between seeing stars from being hit on the head, and seeing points on the ceiling that look like stars: “The difference is that the points look such that one can adjust one's perspective to them” (Textor 2009: 141).

The second premise simply says that if the thesis at stake in P1 is to be given substance, more needs to be said about what aspect of conscious experience makes it so that mind-independence is given in perception. I assume that no one would want to object to this.

The third premise is more contentious, for it says that mind-independence is made manifest precisely by that dual character of perception that I emphasized when first introducing constancy, namely the fact that objects visually appear to retain their properties as the conditions of perception change, while at the same time looking differently as the conditions so change. I think that precisely this duality is required if we are to have a sense of experiencing an objective world in perception. For if things looked to change all the time, we quite obviously could not experience objects as mind-independent. What is perhaps less obvious is that we couldn't perceive them as mind-independent if there were not changes in the ways they appear from different vantage points either. That it is in fact so can be appreciated considering again size constancy. Suppose one sees a tower, and that it looks to be, say, a hundred meters tall. Suppose, then, that one walks some distance away from it, and then looks again at the tower. Wouldn't it be absurd if one were to say that one's senses are deceiving one, because the tower now looks smaller? That this would indeed be absurd is clear when one asks how should the tower look in order for it to look the same size as it did on the first occasion, for the answer certainly cannot be that it should look exactly the same; in that case, one would indeed have grounds for thinking that one is hallucinating, or for supposing there to be some trick in place. It seems, then, that variation in constancy, as it were, helps constituting a sense of mind-independency.

Premise four is the heart of the argument, for it actually connects things as *prima facie* disparate as mind-independence, constancy and the debate over content. What we have up until P3 is only necessary to make sense of perceiving a mind-independent reality, but not sufficient, for we still need an explanation of how this supposed duality of perception actually leads to a coherent picture of the world. What we have is indeed a plurality of facts as to how things look, just as pure relationalists emphasize. But now suppose, as pure relationalists do, that perception did not privilege one of these facts over others. How, then, could we get a coherent picture of reality from perception, instead of a completely schizophrenic representation on which things may or may not change as we move about? Let us consider again the case of the tower seen at a distance. It seems that there are several truths involving the ways it looks: to mention two, it is true that the tower looks the same size as it did when it

was seen close by, and it is also true that it looks smaller. How is it, then, that we typically do not suppose something to be wrong, either with our visual system or in the environment? Or consider again our circular glass seen at a skewed angle: it is true that it looks elliptical, and it is also true that it looks circular. But how, then, do we get a sense of perceiving a mind independent, relatively stationary and unchanging object?

The conception of content as being two-layered gives a straightforward answer to these questions, for it can explain how different facts as to how any one object looks can contribute differently to the constitution of a coherent perceptual content. Thus, it can explain why the glass looks *to be* circular, although it also looks elliptical, and it can explain why the tower looks to be the same size when seen close by as it does when seen at a greater distance, although it also looks smaller on the second occasion. By contrast, it is not clear how pure relationalists could find the resources to answer the questions about the tower and the glass, due to the conception of openness that they embrace – one according to which introducing selection in perception would amount to misconstruing the phenomenological nature of experience. Indeed, for Travis all looks are on a par; but if that is so, it becomes completely obscure how an object can perceptually appear to remain constant in its properties while looking differently under different conditions.

On Brewer's view things are awkward too. Brewer thinks that an object looking F is a question of the object having visually relevant similarities to paradigm exemplars of F-things (I come back to this in chapter three, subsection 4.3.1), where visually relevant similarities include such things as identities in light refraction, retinal images and perceptual processing. Now, in the case of the glass seen at an angle, Brewer's theory gives the wrong result, for it predicts that the glass looks elliptical (from that angle, it has visually relevant similarities to a paradigmatic elliptical thing), but it cannot make sense of the constant aspect of the experience. This strongly suggests that the notion of openness that motivates pure relationalism is unduly narrow and phenomenologically inadequate.

With respect to the first horn of the supposed dilemma for the representationalist, the foregoing argument shows two things: (i) that content is not bound to be incoherent, provided that different looks of an object can play different roles in the constitution of content; (ii) that perception does effect a selection between looks is the best explanation of how we have a sense of being in touch with a mind-independent world in perceptual experience.

With respect to the second horn, a potential objection was that my account implies an arbitrary way of determining perceptual content. In the light of our previous considerations, I think this has bite only when directed against a static, snapshot view of perception and its content. When one considers dynamic experience, one can say that the content of the experience is made manifest to the subject through the whole series of adverbial modifications, and that the content at a given instant is recovered through that series.

In the previous chapter, I have emphasized that intentionality normally involves a perspective on the objects toward which one is directed, and this is literally true for the case of perception. The perspectival nature of perception is of consequence both phenomenologically and epistemologically: concerning the phenomenology, the experience of an object changes as the vantage point changes; concerning the epistemology, certain perspectives are better than others for seeing what properties an object has. The phenomenology of perception is of particular importance for present purposes, because the way in which the phenomenology changes with perspective suggests that there is a two-way explanatory relation between phenomenology and objects: on the one hand, changes in the phenomenology are explained from our point of view as perceiving subjects in terms of mind-independent, relatively stationary and unchanging concrete objects. On the other hand, such objects are revealed to be mind-independent only because the phenomenology changes in particular ways, that is, in ways that are such as to present objects as retaining their features. The conception of content that I've put forward in this chapter is ideal for capturing this structure, for, unlike purely relationalist accounts, it can make sense of objects looking different-ly from different points of view without them looking *to be* different from different points of view. I now address two possible objections, and then conclude with a few remarks on openness and different notions of looks.

2.4 Objections and Responses

We saw a supposed dilemma for the representationalist: either accept that perceptual representation is incoherent, or make an arbitrary choice as to which look indexes content. I consider my proposal a way to avoid both horns of the dilemma. First, we can acknowledge that there are several facts as to how things look, without ending up with incoherent contents, for different looks can play different roles in the determination of content; thus the first horn

is avoided. Second, the determination of the respective function of the looks is not arbitrary, but is phenomenologically determined over the course of a dynamic experience; thus the second horn is avoided as well. There are at least two objections that one might raise: one is a minor point concerning adverbial modifications, while the other presents a more substantial difficulty involving static perception.

a) *Adverbial Modifications*

I have said that there are two layers to perceptual content, and I have also said that one layer represents the object as having a certain property, while the other is an adverbial modification of the former. Intentional content, however, is what is represented *to be* the case. But if the second layer is just an adverbial modification, it doesn't state anything about the way the object is, and so one may wonder why it should contribute to content in the first place. On the other hand, if I were to deny that the second layer is part of the content, I would find myself in the predicament of not being able to assign a definite role to the elliptical look.

I think that the adverbial modifications are indeed part of the content, because they contribute to the constitution of an experience with a certain content rather than another. An experience could not represent, for example, that a glass has a circular rim, if it were not for the elliptical adverbial modifications. If the object were to display, at some point in the dynamic experience, a square apparent shape, the whole content to the effect that the object has a circular rim would be annihilated, because something cannot plausibly look *squarely circular*; therefore, adverbial modifications, in so far as they contribute to the determination of the property that an object is represented *to have*, do constitute a layer of content.

b) *Static Perception*

Even if it is certainly true that dynamic perception plays a fundamental role in our understanding of the world around us, it is also true that sometimes we do perceive statically (we stand still and see stationary objects). How does my proposal fare with respect to these cases? When setting up the dilemma, I mentioned briefly that factors that are strictly speaking non-perceptual, such as expectations, background knowledge, or sensitivity to a particular occasion might have a role in shaping perceptual content and making it accessible to a

subject. The role of expectations is, I think, particularly important in this respect. In a case in which one sees a circular glass, or a surface that is partly in shadow, one has expectations as to how things would look if one were to move about, and those expectations partly determine, together with the ways things appear from the actual point of view, the content of the experience.

There are two features of expectations that make this move coherent with my account. The first is that expectations are not beliefs, and therefore the content of the experience does not turn out to be dependent on the content of beliefs (which would violate Travis's condition that perceptual representation not be autorepresentation). The reason for thinking that expectations are not beliefs is that they do not stand under the same rational requirements as beliefs. For suppose you're looking at a fake barn, while believing it's a real barn instead. If you then were to walk around it and discover it's just a façade, there would be a rational requirement on you to drop your original belief. But if you were to look again at the façade, it might still *seem* to you as if the building actually has a backside, and in that sense there would be an experiential expectation to the effect that the object extends beyond what you can see from your vantage point, where this expectation is not something you could be rationally required to drop.

The second important point about expectations is that they are *grounded* in dynamic perception. Part of what explains you having determinate expectations as to how your experience of a glass would unfold from different perspectives, certainly has something to do with your previous dynamic experiences in which a certain dialectic between the looks of the object has taken place in front of your eyes. Repeated experience sediments to constitute our practical understanding of how things would look if our vantage point were to change. Assigning a role to expectations in the determination and recognition of content, therefore, is consistent both with the primacy of dynamic perception and with the requirement that perceptual content be recognized on the basis of the experience, rather than related beliefs.

Travis thinks that the senses are *silent*, because they bring our surroundings into view, without telling us that they are thus-and-so. His evidence for this claim consists in the supposed lack of principled criteria to determine a coherent content. In fact, the representationalist can admit

that there are several facts as to how an object looks, but still claim that this doesn't jeopardize content. Therefore, she's not obliged to choose between the two horns of the dilemma. She can avoid incoherence because there's a property that an object looks *to have*, as opposed to a merely apparent property. She can also avoid arbitrariness, because which look is the ground look is phenomenologically determined over the course of a dynamic experience; *this* is how perception opens our surroundings to view.

Travis is bound to overlook the distinction between a) the property that an object looks to have and b) the way in which that property looks, precisely because he embraces a distorted conception of openness to the world, on which perception cannot effect any distinction whatsoever between different facts as to how an object looks. But the distinction in question is suggested by the phenomenology of our dynamic experiences and by the way in which we make sense, as perceiving subjects, of the way experience unfolds from different perspectives.

Finally, it is worth noting that my two-layers account is quite neutral with respect to the semantics of look-statements, in that it does not require a special sense or notion of looks in opposition to Travis's demonstrable looks. What we saw with Siegel's objection, was a supposed dichotomy between demonstrable looks on the one hand, and the way things look to a perceiver on the other. What we then saw with Martin's *parsimony* is that there may be no reason to suppose such a dichotomy to obtain. What we can see now is that, even if demonstrable looks were all what one could go on, they do not imply the conception of openness to the world from which *the silence of the senses* originates. For suppose that both the circular and the elliptical look of our glass are facts as to how it looks, in the very same sense of "looks" which that is picked out by Travis's demonstrable looks. Further, suppose that even when we move to partly subjective look-statements, such as "The glass looks circular to S" and "The glass looks elliptical to S", we are still using "looks" in the very same sense as before, as suggested by Martin. Still, this does not mean that the looks cannot play two distinct roles in the determination and recognition of perceptual content. This is entirely consistent with a perfectly reasonable picture of perception as openness to the world, in that (i) both looks are something that belongs to the object, and (ii) their respective role in the constitution of content is determined over the course of a dynamic experience, thus eschewing non-perceptual factors (with the exception of expectations, in the qualified sense analyzed above).

The important result of this chapter is that there is a principled way to determine a coherent, unique, reflectively accessible perceptual content, and that this is compatible with perception as a condition whereby we are directly aware of the world.

Chapter Three: Perception and Illusion

Abstract

In this chapter, I address the question of how one should elucidate perception as openness to the world in the face of illusion and hallucination, and I consider three possible reactions (*sense-data theories*, *representationalism* and *pure relationalism*). In the first section, I show how sense-data theories give up on openness to the world altogether. In the second section, I show how representationalism construes the idea of openness in terms of perceptual representation. In the third section, I expand the dialectical landscape by introducing *naïve realism* and *disjunctivism*. In the fourth section, I consider three purely relationalist accounts of illusion and argue that none of them is entirely satisfactory, although I acknowledge that Brewer's theory is partly successful.

Introduction

If anything is worth the name *openness to the world*, perception surely is. This is firmly rooted in our pre-theoretical conception of experience, in that we all think that perception puts us in touch with concrete reality: it importantly shapes our view of the world, enables us to refer to particular things in our surroundings, and guides our behaviour. At least part of what makes perception so apt at accomplishing these tasks is that perception helps us to get things *right*: it is a reliable source of knowledge in most contexts, and it enables us to cope with potential obstacles in our environment successfully. However, this is not always the case, for sometimes things are different than they appear to be. Thus, in illusion it might look as though something has a certain colour, or shape, or size, or has particular spatial relationships to other things, which in fact has not. In hallucination, it might seem to one as if there is an object with certain characteristics in front of one, although there is no object at all. A significant part of the contemporary debate on perception concerns the question in what sense, in the light of phenomena such as illusion and hallucination, perception can be conceived as openness to the world. Reflection on perceptual error is thus crucial for any theory of perception, not just because of the intrinsically interesting nature of the phenomenon, but also in that different approaches to illusion and hallucination can lead one to different conclusions about the

metaphysical nature of perceptual experience and its objects as such – that is, conclusions that apply generally, regardless of whether a particular experience is veridical, illusory or hallucinatory.

A little more in detail, the dialectic can be set up as follows: on the one hand, we have the intuition that perception puts us in direct contact with reality, that is, that the objects of perceptual experience are material, mind-independent objects out there, and that our experiences depend on them, but not *vice versa*. On the other hand, sometimes things are not quite as they seem to be: in hallucination, it might seem to one that there is a red ball on the table, although there is none; in illusion, it might seem to one that a ball is red, although in fact is blue. Relatedly, we can imagine cases in which one would not be able to discriminate introspectively between a hallucination of a red ball on the one hand, and a veridical perception of a red ball on the other. Introspective indiscriminability is often taken to be of consequence, in that it is thought that for any two perceptual experiences, if they are subjectively indiscriminable, then they are instances of fundamentally the same experiential kind.¹⁴ But since one might think that hallucinations (and possibly illusions) are not actual cases of openness to the world, and since veridical perception, illusion and hallucination all supposedly belong to the same fundamental experiential kind, it becomes difficult to see how veridical perception is openness to the world in the first place.

In what sense, then, could our red ball be the object of our perceptual experiences? Three main interpretations and responses to this question have been put forward over the last century. According to *sense-data theories*, whenever it seems to one that there's a red ball in the surroundings, one actually sees a mind-dependent object in the shape of a red ball, regardless of whether there actually is a material red ball out there. According to *representationalism* (or *intentionalism*, or the *content view*), one represents that there is a red ball, where this requires the existence neither of a mental nor of a material red object, only the occurring of an experience with a certain representational content. Finally, according to *pure relationalism*, perception is essentially a nonrepresentational relation to the object(s) that an

¹⁴ The notion of (in)discriminability that is relevant to the debate on perception has been investigated by Williamson (1990). According to Williamson, discriminating between things can be seen as the “activation of the knowledge that they are distinct, and indiscriminability as the impossibility of activating such knowledge” (1). It is also worth noting that the principle that introspective indiscriminability can be used as a criterion for determining whether two experiences are identical has been subject to severe scrutiny – just as two apples can be qualitatively indiscriminable and yet different objects, some people say, so can two experiences be subjectively indiscriminable and yet different “objects”, or events (for discussion see Williamson (2000, chapter four) and Brewer (2011: 98)).

experience is of, while illusion and hallucination are to be accounted for derivatively with respect to the case of veridical perception, or, as is often called, *the good case*¹⁵.

In this chapter, I want to explore these three approaches. In the first section, I examine sense-data. Sense-data theories have almost disappeared from the contemporary debate, and yet they retain a strong influence, in that the two options which presently receive most attention (representationalism and pure relationalism) can be seen as different reactions to the line of thought which was once supposed to establish sense-data theories. My discussion of sense-data will thus be merely propaedeutic to an understanding of the current debate, and I will skip most of the historical and dialectical subtleties. In the second section, I explain how representationalists resist arguments in favour of sense-data and how they account for perceptual error. In the third section, I expand the dialectical landscape by introducing *naïve realism* and *disjunctivism*. In the fourth section, I examine the purely relationalist account of illusion, with particular attention to Brewer's version, which is arguably the most complete and sophisticated attempt at squaring purely relationalist views with illusion.

1. Sense-Data

The phenomena of hallucination and illusion might, at least *prima facie*, shake our natural conviction that the immediate, direct objects of perception are ordinary, mind-independent things. The reason why it is so, according to a certain line of argument, is the following: when we undergo an illusion we're aware of something that is F, while the external object, by hypothesis, is not F. It logically follows that what we're immediately aware of is not identical with the external object. The conclusion is then extracted that, at least in cases of perceptual error, the immediate objects of perceptual experience are a special kind of mental items (sense-data, that is). The same schema applies of course to hallucination: in hallucination we are aware of an object with certain features, where in fact there is no such object in the external world. It follows that in hallucination we must be aware of some special, mental object, distinct from ordinary objects. Further, it is claimed, since there are no psychologically relevant differences between illusion and hallucination on the one hand, and veridical perception on the other, it must be the case that all what we are ever aware of in a

¹⁵ The good case/bad case terminology comes from Williamson (2000).

perceptually direct, immediate way is just complexes of sense-data. So much for openness to the world!

The relevant arguments might be expounded thus:

Argument from Illusion

P1: In cases of illusion it appears to one that one is immediately aware of a mind-independent object *o* that has a sensible quality F, although in fact *o* is not F;

P2: If it appears to one that one is immediately aware of a mind-independent object that has a sensible quality F, then there is an object that is F of which one is immediately aware (Phenomenal Principle);

P3: The F-object that one is immediately aware of cannot be identical to the non-F mind-independent object that seems to one to be the immediate object of one's awareness (by Leibniz's Law);

P4: Either one is immediately aware of a mind-independent object or one is immediately aware of a mind-dependent object;

C1: In illusion one is immediately aware of a mind-dependent F-object (sense datum/data);

P5: Cases of perceptual illusion can be subjectively indiscriminable from cases of veridical perception;

P6: All subjectively indiscriminable experiences are experiences of the same fundamental kind (Common Kind Assumption);

P7: Perceptual experiences of the same fundamental kind have the same kind of immediate objects;

C2: Both in illusion and veridical perception one is immediately aware of mind-dependent objects (sense datum/data).

Argument from Hallucination

P1: In cases of hallucination it appears to one that one is immediately aware of a mind-independent object *o* that has a sensible quality F, although in fact there is no object at all;

P2: If it appears to one that one is immediately aware of a mind-independent object that has a sensible quality F, then there is an object that is F of which one is immediately aware (Phenomenal Principle);

P3: In cases of hallucination there is no mind-independent object that one could possibly be immediately aware of;

P4: Either one is immediately aware of a mind-independent object or one is immediately aware of a mind-dependent object;

C1: In hallucination one is immediately aware of a mind-dependent object (sense datum/data);

P5: Cases of hallucination can be subjectively indiscriminable from cases of veridical perception;

P6: All subjectively indiscriminable experiences are experiences of the same fundamental kind (Common Kind Assumption);

P7: Perceptual experiences of the same fundamental kind have the same kind of immediate objects;

C2: Both in perception and hallucination one is immediately aware of mind-dependent objects (sense datum/data).

My reconstruction of the arguments, with the articulation in two steps and two conclusions is informed by, although not identical with, that of Snowden (1992).¹⁶ Snowden conceives of the arguments as having a first part that functions as a *base case* for the negative conclusion that in illusion/hallucination we do not directly perceive concrete objects, and a second part that consists in a *spreading step*:

The arguments take a central sort of case, and either claim or argue that in that type of case we do not d-perceive [directly perceive] an external object. I shall call that part of any argument its *Base Case*. Since, however, the arguments are designed to support a completely general conclusion, they incorporate a *spreading step*, which amounts, supposedly, to a justification for thinking that the negative conclusion about the Base Case holds generally of what we, prior to these arguments, take to be perceptual experiences. (Snowdon 1992: 68)

¹⁶ My reconstruction is also similar to that employed by other authors, such as Brewer (2011, chapter 1) Crane (2006, section 3; 2011), Huemer (2011), Robinson (1994: 57-58) and Smith (2002). However, I should note that sometimes the contrast between putative immediate objects of awareness and actual immediate objects of awareness is explained in terms of a difference between ordinary objects and non-ordinary objects, rather than in terms of the mind-independent objects/mind-dependent objects distinction that I employ (see Crane 2011).

The spreading step aims at a generalization from perceptual error to all cases of perceptual experience, based on a phenomenological consideration and a substantive metaphysical thesis. The phenomenological consideration consists in the observation that a hallucination of a red ball could be subjectively indiscriminable from a veridical perception of a red ball. The metaphysical thesis is that subjectively indiscriminable experiences are always experiences of fundamentally the same kind, and hence have objects of the same kind (i.e., sense-data). The metaphysical thesis originates from the significance attributed to indiscriminability, for the intuition that drives P6 and P7 in both arguments is that “there is no non-arbitrary way of distinguishing, from the point of view of the subject of an experience, between the phenomenology of perception and illusions [or hallucination]” (Crane 2011: 2), and that, *therefore*, “there is no reason to suppose that even in the case of genuine perception one is directly or immediately aware of ordinary objects” (2-3).¹⁷

If it is not ordinary, mind-independent objects that one is immediately aware of, then it must be some kind of mental objects, to wit, sense-data. What exactly sense-data are supposed to be is not a question with a straightforward answer. However, Huemer (2011) identifies three conditions that, according to what he calls *the standard conception* of sense-data, would have to be satisfied by an entity in order for it to be a sense-datum:

- i. Sense data are the kind of thing we are directly aware of in perception,
- ii. Sense data are dependent on the mind, and
- iii. Sense data have the properties that perceptually appear to us. (Huemer 2011: 2)

Although sense-data theories used to be quite fashionable in the first half of the 20th century due to the work of authors such as Ayer (1956), Price (1932) and Russell (1912), they no longer enjoy great popularity in the contemporary debate, although see Foster (2000), Jackson (1977), O’Shaughnessy (2003) and Robinson (1994). There are at least two complementary ways in which one might criticize sense-data theories. One is to argue that, given the broad characterization suggested by Huemer, sense-data are a highly problematic kind of entity from a metaphysical, phenomenological and epistemological point of view. A second way is to

¹⁷ The spreading step could also be motivated by an appeal to sameness of proximal neural causes: a hallucination as of a red ball might have the same proximal neural cause as a veridical perception of a red ball. But any two mental events with the same proximal cause, one might think, will be events of the same kind, and thus will have the same kind of object, i.e. a sense-datum. For extensive discussion and criticism of this way of motivating the spreading step, see Martin (2006).

argue that one or more premises in the arguments from illusion and hallucination are false, thus showing the arguments to be unsound. Doing the former would require a lengthy historical discussion of the objections to sense-data and the possible replies to these, and this would bring us quite far from our central concern with perception as openness to the world.¹⁸ I will rather do the latter, which will take us directly to the heart of the disagreement between representational and purely relational approaches.

2. Representationalism

The focus of our enquiry concerns the question of what metaphysical consequence (if any) has to be drawn from the analysis of illusion and hallucination, in light of the possible subjective indiscriminability between falsidical and veridical experiences. These supposed consequences are twofold: they might concern the nature of the objects of perception, or the nature of perception as a mental state. We've seen in the previous section that, according to sense-data theorists, the consequence to be drawn is that hold that the immediate objects of perception are always mind-dependent.

In this section, I examine the representationalist story, and explain how representationalists think that there is a conclusion concerning perception as a mental state to be drawn from the argument of illusion/hallucination. Other than expository purposes, this section also has two critical aims: first, I want to argue that if representationalists want to give a satisfactory account of perception as openness to the world, they have to incorporate the relationalist insight that perceptual experiences should be individuated partly in terms of their objects. Second, I defend the representational view from two criticisms put forward by Brewer.

Let us consider again the argument from hallucination:

¹⁸ For a general overview of these objections, see Crane (2011, section 3) and Huemer (2011). More in particular, metaphysical objections against sense-data are raised by Armstrong (1968). A phenomenological objection is raised by Husserl (2001: 125), who argues that supposing sense-data to be similar to mind-independent objects, as is normally done by sense-data theorists, is a phenomenological nonsense; the same line is taken up by Burge (1986: 128) and Searle (1983: 59-60). Familiar epistemological worries are discussed by Huemer (2011, section 3.2).

Argument from Hallucination

P1: In cases of hallucination it appears to one that one is immediately aware of a mind-independent object *o* that has a sensible quality *F*, although in fact there is no object at all;

P2: If it appears to one that one is immediately aware of a mind-independent object that has a sensible quality *F*, then there is an object that is *F* of which one is immediately aware (Phenomenal Principle);

P3: In cases of hallucination there is no mind-independent object that one could possibly be immediately aware of;

P4: Either one is immediately aware of a mind-independent object or one is immediately aware of a mind-dependent object;

C1: In hallucination one is immediately aware of a mind-dependent object (sense datum/data);

P5: Cases of hallucination can be subjectively indiscriminable from cases of veridical perception;

P6: All subjectively indiscriminable experiences are experiences of the same fundamental kind (Common Kind Assumption);

P7: Perceptual experiences of the same fundamental kind have the same kind of immediate objects;

C2: Both in perception and hallucination one is immediately aware of mind-dependent objects (sense datum/data).

A representationalist would typically resist the argument by denying P2 both in the argument from hallucination and in the argument from illusion. P2 is actually a particular way of expressing a general idea that goes under the name of *Phenomenal Principle*. Robinson (1994) formulates it thus:

If there sensibly appears to a subject to be something which possesses a particular sensible quality then there is something of which the subject is aware which does possess that sensible quality. (Robinson 1994: 32)

Focusing on the case of vision, we can reformulate the principle in terms of the way things look to a subject:

If something looks F to a subject S, then there's something of which the subject is aware which has the property F.

However intuitive this may seem at first, it certainly loses all its appeal when we apply it to intentional states such as belief or judgement. For consider the following:

If a subject judges (desires, fears, believes...) that an object has the property F, then there's something which has the property F.

The latter principle is obviously false, and a representationalist would typically say the same of the former. If I desire to ride a black dragon, this does not mean that there must actually be an object which is a black dragon, and if I believe that there's a shop at the corner of the street, it doesn't mean that there must be a shop there. Likewise for perceptual experience: if I enjoy an experience as of a red ball, that does not imply that there must be a red ball (either mental or physical) of which I'm aware; what is required is simply a representational content to the effect that there is a red ball.

The contrast between the phenomenal principle and the representational approach brings out a metaphysical assumption that underlies sense-data theories – one which is only implicitly present in the argument from hallucination/illusion: that perceptual experience is essentially relational. *One* kind of representationalist approach precisely denies this relational character. According to that kind of representationalism, the denial of the phenomenal principle is based on the denial of a more general assumption about experience: that it is a relation to a real object. Rather, this sort of representationalist would say, perceptual experiences are intentional states which represent the world. Thus Crane says:

Perception is a representation of the world; and... it's not generally true that if X represents Y, then, Y must exist. So a perceptual experience need not essentially involve a relation to what it represents. This is the intentionalist conception of perception. (Crane 2006: 133)

This is not exempt from problems, though, because intentionality, and specifically perceptual intentionality, is supposed to be the property of our mental states of being about objects in the

world. What Crane says, however, is that X can be about Y, regardless of the existence and features of Y. This seems to generate a tension: on the one hand, X is about Y; on the other hand, it makes no difference whatsoever to X whether or not Y exists or not. Crane is aware of this, but accepts with equanimity that “perception ‘falls short’ of the world” (Crane 2006: 141); however, if one wants to explicate perceptual consciousness as openness to the world, one will not be at ease with this outcome.

At the level of logical form, a relevant distinction to be drawn here is that between existential and singular contents. Suppose that one sees a red sphere. The contrast is as follows:

Existential: there is an x, such that Red(x) and Spherical(x);

Singular: Red(a) & Spherical (a).

If the contents of perception are existential, all what is “said” by a perceptual experience as of a red ball is that there is *an* object which is red and spherical, without specifying which object is putatively being perceived. The truth-conditions that existential contents determine are satisfied by whatever object in an appropriately demarcated spatial region satisfies the general predicates *being red* and *being spherical*. If, by contrast, the contents of perception are singular, then perception purports to make reference to one particular object. Singular contents are object-dependent contents: the satisfaction of singular contents depends on the same individual across possible worlds, not the presence of whatever satisfies, in a purely descriptive way, the general relevant predicates. Already in chapter one, (2.2) I committed myself to the existence of singular, *de re* contents for some thoughts and other intentional states, and I identified in a contextually-anchored, demonstrative perceptual reference the relational basis for those contents. Back then I did not spend many words explaining why one should think of the content of perception as singular, though. This seems the appropriate moment to do this, as the challenge for the representationalist is to explain how in cases of illusion there still is a particular mind-independent object that one sees – to explain, that is, why on a representationalist view perception does *not* fall short of the world, *pace* Crane.

There are good reasons for thinking that the contents of perception are singular, and I think most representationalists would agree, although see Davies (1992) and McGinn (1982). A powerful argument for the object-dependency of perception has been offered by Soteriou

(2000)¹⁹, who argues that the question of the veridicality of an experience cannot be “settled independently of the question of which particular object is being perceived” (179). The argument proceeds by means of a few intuitive steps, and effectively shows how a supporter of general contents is problematically committed to the possibility of *veridical misperception*.

First, consider a scenario in which a subject thinks she sees a red sphere just in front of her, but she is unknowingly wearing distorting glasses that shift everything to the right, so the object is in fact to her left. Intuitively, the subject is misperceiving the object, and her experience is non-veridical.

Second, suppose that we place a red sphere just in front of the subject, where she thinks that a red sphere has been all along. Due to the distorting glasses she is wearing, the subject has the impression that another red sphere has been added to the right.

Third, suppose we add yet another red sphere to the right – “an object that may be too far to the right of the subject to be perceived by her” (Soteriou 2000: 180).

Fourth, a plausible specification of the content of the experience in general terms would be something like: “there is an object x which is red, spherical and in front of me, and there is an object y that is red, spherical and to my right”. Importantly, this content would be *completely veridical*, for it does not say that there is no object to the left (which would be false), and it is actually true that there is a red sphere in front of the subject and another to the right. At the same time, the experience would also be a case of misperception, for the object which the subject thinks is in front of her is in fact to the left, and the object which she thinks is to the right is in front of her; therefore she misperceives two objects. Thus the supporter of general contents should say that the experience in question would be a case of *veridical misperception*, because the content conceived of in general terms is completely veridical, but at the same time the subject misperceives her environment.

Fifth and final, the idea of veridical misperception is a monster that should not be allowed to enter one’s story of perception, for it implies the denial of the following plausible principle:

¹⁹ There is at least another, *epistemological* argument that has been put forward by Campbell (2002a) for the related claim that perceptual experiences should be partly individuated in terms of their objects: it seems that perception puts us in a position to formulate knowledgeable demonstrative judgements about things in our environment. Now, this implies, says Campbell, that we have knowledge of the reference of the demonstrative terms we employ in making the judgements in question. Such knowledge, however, can only be available to us if perceptual experiences are individuated partly in terms of the particular objects that we see. I should also mention that there are phenomenological and metaphysical themes that are relevant to the particularity of perception, for it is claimed by some that it is the particular object that we perceive which constitutes the phenomenal character of experience, and that that object is indeed a constituent of the experience (Martin 1997, Fish 2009, Logue 2013); more on this in section three of this chapter.

“if some part of the subject’s environment is different from the way it is represented to be, then at least one of the conditions required for the content of the representation to be correct is not satisfied” (180). However, it seems clear that we should not deny the principle in question:

We would be loath to give up an analogous assumption in the case of belief. If we know that some part of the subject’s environment is believed by her to be a certain way, and we also know that that part of the subject’s environment is different from the way that it is believed by her to be, we feel we have sufficient grounds to claim that the subject has a false belief. (180)

It is especially relevant here that it is an analogy with other intentional states that tells against the hypothesis of existential contents, for the very way in which representationalists motivate their denial of the phenomenal principle implies very significant commonalities between perception and belief – i.e., just as believing something to be F does not imply that something is F, so also an event of it looking to one that something is F does not imply that something is F. Likewise, just as believing the environment to be other than it is implies that the content of one’s belief is false, so also perceiving the environment to be other than it is should be taken to imply that one’s perception is non-veridical.

The steps in Soteriou’s argument seem unassailable, so representationalists should incorporate an element of relationality in their account, in the shape of singular, *de re* contents. Advocates of singular perceptual contents have no difficulty agreeing with other representationalists *à la* Crane over the denial of the phenomenal principle, for they are by no means compelled to say that if it seems to one that one sees that an object is F, then they are related to some object that is F. However, unlike Crane, they can accommodate perception as openness to the world, since the particular object that one sees on a given occasion is not something extraneous to the experience that supposedly concerns it. Thus in accordance with the first critical aim of this section, I contend that a version of representationalism which endorses singular contents appears to avoid spooky entities such as sense-data and dubious theoretical assumptions like the phenomenal principle on the one hand, while it can account for openness to the world better much better than sense-data theories on the other.

The second critical aim of this section is to defend representationalism from two criticisms put forward by Brewer. The first criticism is an objection to the intentionalist treatment of a particular case of illusion, and actually consists of two challenges.

Consider the Müller-Lyer diagram (ML), in which two lines look different in length, while in fact they are exactly the same length, and the impression of inequality is due to the fact that one line has inward fashes and the other has outward fashes. Typically, a representationalist would say that a perceptual experience of the diagram falsely represents the lines as being different in length; the experience concerns that very diagram, constituted of two equal lines, but they are represented as being unequal. Brewer highlights two difficulties for the representationalist treatment of ML, which he thinks are symptomatic of a more general problem – one which originates from a theoretical commitment essential to representationalism, namely that, for any given experience, there must be a specific way the world would have to be in order for the experience to be accurate. Brewer thinks that ML brings into relief problematic aspects of this commitment. Here is his first challenge, in the guise of a series of questions:

...is the line with inward fashes supposed to be represented as shorter than it actually is;
or is the line with outward fashes supposed to be represented as longer than it actually is;
or both; and by how much in each case? (Brewer 2011: 65)

As was the case with Travis's argument that we saw in chapter two, the challenge is to select in a principled way one from among many possible perceptual contents, all of which would be compatible with the way things are, as *the* representational content of a given experience. In this particular case though, the worry could be avoided precisely by *not* answering the questions that Brewer asks. Thus, McDowell notices:

Suppose I say of two lines that are in fact the same length that one, say A, is longer than the other, say B. In saying that, I am representing A as longer than B. it does not follow that I am saying that A is longer than it is, or saying that B is shorter than it is, or saying that both these things are the case. One of these things would have to be so if what I say were true, but I am not saying of any one of them that it is so. A "by how much" question does not arise. Just so with an experience that represents one line as longer than the other. (McDowell 2008: 201)

I think this is an adequate way of responding to the challenge, except for one thing: the claim that a “by how much question”, in thought as much as in perception, does not arise. As Brewer points out, I can surely think or say that “A is longer than B without saying anything at all about how long either A or B is, or by how much” whereas in the case of seeing “I normally see that A is longer than B by seeing the extent of A and the extent of B and noticing the former is greater than the latter by roughly such and such an amount” (Brewer 2011: 66). The point applies generally, of course: I can think that there is a cup on the table, period. By contrast, I cannot see that there’s a cup on the table, without any detail concerning the shape, colour or location of the cup figuring in the experience. So the worry is that seeing involves a far greater degree of determinacy – it displays more *fineness of grain*, as they say²⁰ – than thought, and it seems that if one takes a McDowellian stance by saying that in ML line A is represented as merely determinably longer than B, then one is neglecting the obvious phenomenological fact that perception is rich in detail. I think that Brewer is right on this.

However, I also think that there is an easy way out of the problem. The representationalist not being able to say *exactly* by how much A is represented as longer than B doesn’t amount to a denial of the fineness of grain of experience, for fineness of grain certainly does not imply perfect acuity. One might then adopt this position: the content of an ML-experience is one to the effect that A is longer than B – which doesn’t require a specification on whether A is represented as longer than it is, or rather B as shorter than it is. There is, however, a certain “by how much” question, in the sense that it’s not indeterminate whether A is represented as being a couple of meters or a couple of centimeters longer than B. Although unable to specify a precise length, I would say that A looks *about* one centimeter longer than B (but this obviously depends on the size of the particular diagram that one is looking at) – thus a more full specification of the content of an ML-experience might be *A is longer than B by this much*, where “this much” corresponds to the perceived difference of about one centimeter. Therefore, the representationalist is neither obliged to give a random answer (thus reinforcing the relationalist conviction that there’s no principled criterion to individuate the content of an experience) nor is she obliged to deny that the question arises in the first place (which would be phenomenologically off-key). Instead, she can give an answer which respects the fine

²⁰ The fineness of grain of perceptual experience has also been taken as evidence to the effect that perceptual content is nonconceptual by Evans (1982), among others. For a conceptualist response see McDowell (1994, Lecture III).

grained nature of experience, although of course one which falls short of perfect acuity. But since fineness of grain does not require perfect acuity, this is unproblematic.

Brewer's second remark has the aim of bringing out a tension between perception as an encounter with the actual constituents of reality on the one hand, and a logical impossibility that would have to figure in the representational content of an ML-experience on the other hand:

...[Representationalism] appears committed to some kind of conflict within the content of the (ML)-experience itself. For the endpoint of each line certainly looks to me to be exactly where it actually is, at its actual position on the (ML) diagram...At the same time, (CV) [Representationalism] registers the illusory nature of my experience with the claim that it is also part of its representational content that A is longer than B...(CV) is thereby committed to regarding the representational content of my (ML)-experience as *impossible*: it cannot possibly be veridical. (Brewer 2011: 66-67)

I think there are a number of things that can be said in response to this. The simplest thing is to challenge the claim that we actually represent the endpoints of the main lines accurately. And if Brewer were to appeal to action to back up his point, for example saying that we would accurately reach for the ends of the lines or point to them, then this strategy can be argued against, by highlighting that representations for action are different from those of visual experience – that if given the task of reaching out for the endpoints we would carry it out successfully, does not entail that those endpoints are visually represented accurately. But if we don't represent the endpoints accurately in the first place, then any possible problem deriving from impossible contents is blocked right from the start.

A second response would make appeal to demonstrative modes of presentation. The representationalist might accept that the subject represents the endpoints correctly as being there (a), there (b), there (c) and there (d). Suppose also that the lengths are represented as "that long (A)" and "that long (B)". Then there is no straightforward impossibility in the content, as the equality of the lines is not explicit in the modes of presentation, only implicit in their referents.

A third response would consist in accepting without feeling inconvenienced the possibility of perceptual representational contents that could not possibly be true. One might indeed say that impossible contents are something that speaks in favour of the idea that

perception has nonconceptual content, not in favour of the claim that perception does not have content at all.²¹

On balance, it seems that the local objections against the representationalist treatment of ML are not decisive.

Brewer also raises a second, more fundamental problem for representationalism, which concerns a tension between the unboundedness of perceptual content, as it were, and the limits beyond which an object simply fails to be presented in perceptual experience, regardless of its causal role in triggering a mental episode in a subject. Upon seeing an object *o*, there seems to be no bar to how wrong a representational content concerning *o* might be. Supposing *o* to be a red ball, all of the following might be, by representationalist lights, perceptual contents concerning it: <*o* is a green ball>, <*o* is a white glass>, <*o* is a perfect square>, <*o* is a pink elephant>, and so on. However, there is *some* limit beyond which the object in question fails to genuinely figure in one's experience, in the sense that it is not *seen*, even if it plays a role in causing the relevant mental episode. Which point exactly this may be is not a question with an easy and definite answer, for it may largely depend on contingent, empirical factors that are not identifiable through conceptual analysis. But the representationalist, or so Brewer's objection goes, is not even in a position to start to give an answer, for there is nothing in the notion of representational content as such that may impose a constraint, even a very approximate one, for determining when a given representational episode counts as a seeing, and when it doesn't because the predicative part of the content ascribes to the putatively seen object something that is just miles off the truth. In Brewer's own terms:

One the one hand, the phenomenology of genuine perceptual presentation surely places certain limits on the nature and extent of any errors involved. On the other hand, the basic notion of false content...appears subject to far less demanding, if any, such limits.
(Brewer 2011: 71)

One way of appreciating the nature of the challenge is to consider the difference between an episode of *seeing* an object *o* on the one hand, and an act of *referring* to *o* on the other. Representationalism construes perceptual experience in terms of a representational content

²¹ A similar line is put forward by Crane (1988).

that concerns, or purports to refer to, some particular object *o*. This means that, according to representationalism, the limits beyond which an experience fails to concern *o* are the limits beyond which reference fails. Just what such limits are is of course a large a question which I am not in a position to address here, but all what I need to emphasize for present purposes is that reference is a quite tolerant relation, in that acts of genuine reference are compatible with wild error in the attribution of properties to the thing that is referred to. Thus, if I say that the man standing at the bar is the boss of the local criminal gang, whereas he is in fact the meekest guy in town, I am very mistaken about the nature of the man, but I still succeed in referring (perhaps demonstratively) to him. The relation of *seeing*, on the other hand, is less tolerant. Intuitively, if *o* is a red ball, and I point to it and sincerely say that it looks like a pink elephant, then I fail to see *o*. Brewer explicitly admits that these sought-for limits on perceptual presentation, or more specifically on seeing, are “not hard and fast”, and he adds: “there may be wildly abnormal circumstances in which the illusions I reject are in fact possible” (73). Yet, he insists that “given a fixed environment and evolutionary context, there are limits beyond which an object fails to be genuinely presented in perception” (74).

The present point has been one of the subjects of a recent exchange between Adam Pautz and Brewer, with Pautz arguing that problems concerning the definition of perceptual presentation are anyone’s problems:

If you try to take a picture of your cat, but your camera is very (very!) broken so that what comes out happens to look like Eiffel tower, then we don’t count the picture as a picture of your cat. The explanation is that it is a just a fact that the picture of relation requires a suitable degree of match and a causal process that is not too deviant. Likewise, I think that the content theorist should say that it is just a fact about seeing that, if you are to count as seeing an object, then that object’s properties must suitably match the properties attributed by the content of your experience. (Pautz, unpublished: 3)

Analogous things, according to Pautz, could be said for knowledge, or beauty: there are no hard and fast limits on what kind or amount of luck is compatible with knowledge, or on which features of an object are compatible with its beauty. With respect to knowledge, Pautz elucidates his point by means of an analogy with a Williamsonian account:

Timothy Williamson holds that believing is a common factor between knowing and not knowing. He also holds that truth and sensitivity are necessary conditions on knowledge.

But he doesn't hold that knowing is a composite condition built up from belief, truth, sensitivity, and other factors. He advocates a non-composite view of knowing, because of the dismal history of conceptual analysis. The content theorist could take an analogous, Williamsonian non-composite view of seeing, but still say that causation helps ground facts about what you see and what you don't see. (5)

So Pautz's idea seems this: just as knowledge is primitive but being in a state of knowledge implies that other conditions obtain, so seeing too is primitive (a non-composite condition) but an episode of seeing implies that other conditions such as causal factors obtain, and these causal factors partly determine what one does or does not see.

Brewer's response is that it matters where one starts, that is, it makes a difference whether one takes the notion of perceptual relation to an object as primitive, thus derivatively explaining facts as to how that object looks, or rather the notion of representational content as the fundamental *explanans*:

Conscious acquaintance, actually seeing o, is the most fundamental condition in the area. Taking this as basic it is possible to explain the various ways that things look in perception. Taking the way things look, which may be either veridical or illusory, as basic, on the other hand, as (CV) [representationalism] does, it is impossible to work up from there to an adequate account of seeing the things in the world around us (Brewer 2015: 2.1).

Further, Brewer thinks that his brand of relationalism does in fact enable him to give a principled, if approximate, answer to the question of the limits beyond which an object fails to be presented in perception:

(OV) [relationalism] does have an explanation of the limits of error in illusion compatible with seeing, although of course, it starts with seeing itself and derives the limits on error rather than vice versa. Given acquaintance with o, there are limits to the way o looks: just those ways, E, such that o has, from the point of view of acquaintance and in the circumstances in question, visually relevant similarities with paradigm exemplars of E (Brewer 2015: 2.1).

The dialectic is quite complex here. On the one hand, one may find (as Brewer does) infelicitous the analogy that Pautz draws with a Williamsonian account, for Williamson's

point is precisely that knowledge is explanatorily prior, and that concepts such as that of justification can be illuminated by it, not *vice versa*; however, one might say, it is not obvious that a friend of content can embrace this line for cases of seeing. According to a representationalist approach, a perceptual experience of an object *o* consists in perceptually representing that *o* is some way. But then, Brewer's line goes, the representationalist cannot have a principled criterion to determine whether or not a given representational episode can count as one of seeing – the representationalist is unable to reconstruct the notion of seeing in terms of representational and causal facts.

On the other hand, things are not straightforward for Brewer either, although it takes a few steps to see why.

First, let us consider the question: is it true that if one is perceptually acquainted with a particular object, then that object will look some way to one? I would say it is, or else understanding the fundamental role of acquaintance in the elucidation of the phenomenology of perception would become difficult. Indeed Brewer accepts this, for he writes that “if a subject is aware of the world, in vision, say, then there are truths of the form “*o* looks *F*” that qualify in a perfectly natural sense to capture how things seem to the subject” (Brewer 2011: 62).

Second, Brewer thinks that the way(s) any one given object looks to a subject that is acquainted with it, is a function of the visually relevant similarities (see 4.3.1 below) that, given particular circumstances, the object has to other kinds of objects.

Third, suppose scientists were to discover that, under very unusual circumstances, there are visually relevant similarities between a pair of equal lines and a perfect circle.

Fourth and final, suppose that one stood in front of a pair of unequal lines, with eyes open and functioning, in a situation in which the lines do have visually relevant similarities to a perfect circle, and that one had a visual experience as of a perfect circle.

Now, it seems to me that Brewer would face a difficult choice. He could say that the subject in question does see the two lines, although they look like a circle to her. However, this would amount to accepting a reasoning from causal factors and factors about looks to acquaintance – that is, it would be to use facts about looks in order to determine whether or not one sees an object. This should not be very welcome by Brewer, for the charge that he moves to representationalism is precisely that of not being able to accommodate the primitive

nature of perceptual presentation due to a mistaken explanatory emphasis on the way objects look.

Alternatively, he could stick to his guns and say that the subject fails to see the lines, regardless of their causal role in triggering the relevant mental episode. This raises an awkward question, though: what is it, if not the two lines, that looks like a perfect circle to the subject? Brewer would have to say that nothing does, for the subject is merely hallucinating. While this would be perfectly coherent with the general framework that he endorses, I find it counterintuitive, for we would have a situation in which (i) *o* does have visually relevant similarities to a circle, (ii) *o* and its visually relevant similarities to a circle play an appropriate causal role in triggering a corresponding experience, (iii) the subject looks in the right direction, with a sufficient degree of attention, but still (iv) she fails to see *o*. This, it seems to me, would be a dogmatic way of overlooking facts about causation which could help determining what the subject does or does not see.

If my reasoning is correct, Brewer and the representationalist are pretty much in the same situation. The representationalist may not have a clear-cut criterion or even a guiding principle for determining the point beyond which a given representational episode stops being a seeing. Brewer contends that he is in a position to draw such a distinction, in terms of the visually relevant similarities that an object of acquaintance has to other objects – that is, given that one sees *o* in the first place, then the limits to the ways that *o* looks are determined by its visually relevant similarities to several kinds of objects. However, the case of the hypothetical pair of lines which looks like a circle suggests that one might start with facts about looks, together with other causal factors, and then move to the conclusion that one does (or does not) see the two lines in question. In other words, it seems that the explanatory direction *could* go from facts other than seeing itself to seeing, rather than *vice versa*. On balance, it does not seem that Brewer's objection is decisive, or that he can show his brand of relationalism to have an overwhelming explanatory advantage over representationalism.

Pure relationalists also have positive things to say about illusion and perceptual experience in general. In the next section I discuss their position in more detail, especially in connection with two related ideas: *naïve realism* and *disjunctivism*. In the fourth section, I offer a detailed discussion of the purely relationalist account of illusion.

3. Pure Relationalism, Naïve Realism and Disjunctivism

Over the course of chapter one and two, I have expounded pure relationalism as a view that results from a particular theoretical gloss on perception as openness to the world, conceived of in terms that are such to rule out perceptual content. One aspect of openness to the world as conceived of by pure relationalists is that perception does not select one from among many facts as to how any one object looks – that was the main theme of chapter two and the point of Travis’s argument. Other two, related aspects will be under our lens in this chapter: the passivity of perception, and the role of objects in the elucidation of perceptual experience.

A purely relationalist insight is that if perception really is to be an encounter with the layout of reality, then it has to be passive, and as such completely innocent in connection with the possibility of a subject getting things wrong in the light of her experience. In this connection, contemporary relationalists often trace the origin of their ideas back to the British Empiricists although, as remarked by Travis (2004), even Descartes seemed to share some of their basic insights. Thus Descartes writes, on mental states in general:

By the mere intellect I do no more than perceive the ideas that are matters for my judgment; and precisely so regarded the intellect contains, properly speaking, no error. (Descartes 1971: 99)

Whence then, do my errors originate? Surely just from this: my will extends more widely than my understanding, and yet I do not restrain it within the same bounds, but apply it to what I do not understand. (Descartes 1971: 96-97)

A well-known passage from Hume also contains the claim that the relation between the basic materials offered to the mind in its encounter with the world (impressions) and their conceptual counterparts (ideas) is one of *exact* representation:

...all our simple ideas in their first appearance are deriv’d from simple impressions, which are correspondent to them, and which they exactly represent. (Hume 1960: 4)

One should not be led astray by the word “representation” here, for it means a different thing altogether than it does in the mouth of contemporary representationalists. When introducing the notion of intentionality in chapter one, I said that intentional states represent their objects

in a certain way. When applied to perception, this means that there is a way things look for the subject in having an experience, where the way things look indexes the representational content of the experience, which may or may not reflect the way things actually are. But this is not how Hume uses “representation”. Hume rather means a pictorial relation of similarity like one that might hold between, say, an apple and a picture of an apple. If we then add the requirement that similarity be *exact*, we obtain a picture of perception that is importantly similar to that endorsed by contemporary pure relationalists: one on which (i) perceptual consciousness is to be understood in terms of a relation of acquaintance with its direct objects (impressions for Hume; mind-independent objects for contemporary relationalists), and (ii) representational error, strictly speaking, may only occur in thought.

In the last century, the theme of the passivity and epistemic innocence of perception has been taken up by Austin (1962), who complained that talking of perceptual experience as involving a commitment as to how things are, and so as liable to be inaccurate, could be a metaphor at best, which leads one astray as soon as it is taken literally:

...Though the phrase “deceived by our senses” is a common metaphor, it *is* a metaphor; and this is worth noting, for...the same metaphor is frequently taken up by the expression “veridical” and taken very seriously. In fact, of course, our senses are dumb – though Descartes and others speak of “the testimony of the senses”, our senses do not *tell* us anything, true or false. (Austin 1962: 11)

This is also very much emphasised in the most recent formulations of purely relationalist approaches to perception. Gupta, for example, thinks that perceiving is a bit like bumping into a tree:

...experience is passive, and it is always a good policy not to assign fault to the passive. If during a walk in a forest I bump my head on a low branch of a tree, it is better that I assume responsibility (and change my ways) than that I pin the blame on the tree. The tree is passive. It is bound to be the way it is, given the circumstances; and it is useless to blame it for my sore head. Similarly, if having suffered an experience, I acquire a false perceptual belief, it is better that I assume responsibility (and change my manner of “reading” experience) than that I pin blame on the experience. The experience is bound to be the way it is, given the circumstances; and it is useless to blame it for my false belief. (Gupta 2006: 185)

And thus Travis and Brewer, in a fashion that will be familiar by now, say:

...*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 2004: 65)

...the possibility of falsity is a net cost, not a benefit, to (CV) [representationalism]. (Brewer 2006: 172)

In addition to the passivity of perception, contemporary pure relationalists emphasise, just like the British Empiricists did, that perception is a condition that is to be elucidated with reference to its direct objects — for contemporary relationalists, these are the concrete particulars we encounter in everyday life. The analysis of this particular aspect of perception as openness to the world and its implications for the topic of perceptual error requires an expansion of our dialectical landscape, with the introduction of *naïve realism* and *disjunctivism*. I should also emphasise that, while pure relationalism certainly entails naïve realism and disjunctivism, it is not obvious that one could not be a naïve realist and a disjunctivist without being a pure relationalist; that is, it might be that one could be a naïve realist and a disjunctivist while at the same time endorsing perceptual content — this section is mainly expository, so I avoid this issue here, but I come back to it in the concluding remarks of the thesis.

Naïve realism (NR) has become increasingly popular over the last decade, and is meant by its proponents to be the best theoretical articulation of “how our experiences strike us as being to introspective reflection on them” (Martin 2004: 42). So naïve realists intend to accommodate the most important traits of what they take to be a pre-theoretical view of perception, and in this sense they see themselves as “saving appearances”, as it were, whereas sense-data theories and representationalism constitute, according to Martin “error-theories of sense experience” (Martin 2004: 84). This difference of approach between sense-data theories and representationalism on the one hand, and NR on the other, does indeed emerge clearly in the treatment of perceptual error which they offer.

Sense-data theories, and to a lesser degree representationalism, are both motivated by the argument of illusion/hallucination, and thus their proponents typically accept that one should move from considerations pertaining perceptual error to conclusions about the nature of perceptual experience in general: the nature of perceptual failure dictates what theory one

should embrace of perceptual success.²² This certainly is a methodologically sensible approach. If you want to know how a machine, process or mechanism works, it might be a good idea to try to remove one component or another and observe how the thing fails to accomplish its function as a consequence; the study of failure is explanatorily valuable in the understanding of success. Likewise for perception: the study of hallucination and illusion is explanatorily valuable for the understanding of the nature of perceptual experience in general.

This marks out sense-data theories, as well as representationalism (but only insofar as it is motivated by the argument of illusion/hallucination), as what may be called *scientific-like* approaches to perception, in the sense that they start out with the examination of a theoretical problem and derive consequences for perception in general. The particular form that their inquiry take is akin to that of vision sciences. One of the major ongoing research projects in vision sciences is that of figuring out how our visual system can get to a coherent and determinate picture of the environment, given the relative poverty of the sensory input that it receives, and given that the two-dimensional image projected onto our retina is compatible with an *infinite*²³ number of three-dimensional arrays of reality. A very influential response to this question is that there must be certain *assumptions* hard-wired into our visual system, that link variations in light distribution to elements in a three-dimensional array, such as margins, corners and colours. It is in virtue of these assumptions that the visual system can discard implausible scenarios, thus making it so that most of the time we can cope with our environment successfully; it is also because of these very same assumptions that sometimes our visual system falls prey to illusions. Thus vision sciences start out considering a certain explanatory problem in the form “how, given X, is it possible that Y?”, and proceed to give an answer that is of value for a general understanding of the way our visual system works.

Sense-data theories and representationalism approach the question of perception in a way that has important similarities to that of vision sciences, in that they investigate a question of the form “how, given the possibility of illusion and hallucination, can perception be openness to the world?”. The fact that they give very different responses to the question (sense-data theories denying that perception is openness to the world in the first place, and

²² The argument from illusion/hallucination is not the main reason why I endorse representationalism (my arguments are more phenomenologically-oriented), but there certainly are authors for whom this is a major motivation (see for example Crane 2001, 2006).

²³ This is not an overstatement; for discussion see Palmer (1999:23).

representationalism construing the notion of openness to the world in terms of representational content) should not make us lose sight of a fundamental common ground that those theories share, namely the assumption that whatever answer one may want to give to the question of openness to the world, it will have to be shaped by an examination of illusion and hallucination.

By contrast, *NR starts out as a theoretical elucidation of a pre-theoretical view of veridical perception*²⁴, and is as such silent on perceptual error. NR holds that veridical perceptual experiences have material objects as their *constituents*. This implies that two experiences of two numerically different, although qualitatively identical objects cannot be instances of fundamentally the same experience, for experiences are individuated partly in terms of the experienced objects. An influential proponent of NR such as Martin writes:

Some of the objects of perception – the concrete individuals, their properties, the events these partake in – are constituents of the experience. No experience like this, no experience of fundamentally the same kind, could have occurred had no appropriate candidate for awareness existed. (Martin 2004: 39)

Heather Logue follows in the same footsteps, although she offers a different formulation of the idea of an object being a constituent of perceptual experience:

According to my characterization, a veridical experience fundamentally consists in the following state of affairs: the obtaining of the perceptual relation between the subject and the mind-independent objects of the experience. Thus, any constituents of this state of affairs are constituents of the experience (including the entities perceived). It's not clear how else to understand the idea of a mind-independent object being a constituent of an experience other than in terms of the experience being a state of affairs that has the object as a constituent. (Logue 2013: 107, fn. 5)

Sometimes the point that perceptual experience is to be elucidated with reference to its objects is expressed with an emphasis on the subjective character of experience. Thus Brewer writes:

²⁴ A notable exception is Logue (2012); despite being a naïve realist, Logue does not think there is such thing as “the pre-theoretical view of veridical perception” (not in a sense that would be relevant to the debate over the metaphysics of perceptual experience, at any rate). Of course, if she is right, the label *naïve* realism is an inapt one.

The idea that the core subjective character of perceptual experience is to be given simply by citing the object presented in that experience is more fundamental than any appeal to perceptual content. (Brewer 2008: 168)

What emerges from these passages is that there are two complementary ways of elucidating perception as a conscious relation of acquaintance with particular objects from a naïve realist perspective: one with explicit reference to the metaphysical composition of the experience and one with reference to the elucidation, or, in more theoretically loaded terms, “constitution” (Fish 2009: 6), of the phenomenology of experience in terms of concrete, mind-independent particulars.²⁵ Both formulations face very similar challenges. For suppose one said that material objects are constituents of veridical perceptual experiences. Then, one obviously would have to give an alternative explanation of the metaphysical composition of hallucinatory (and possibly illusory) experiences, since there is just no object there in hallucination. Likewise for phenomenal character: if objects constitute the phenomenal character of perception, an alternative explanation must be given for the phenomenology of hallucination.²⁶ Since NR is just really about veridical perception, it is not equipped, if taken by itself, to answer questions concerning illusion or hallucination. However, naïve realists can help themselves to a related idea: *disjunctivism*.

Disjunctivists start from a very simple observation: if it perceptually seems to S that *p*, then it might *either* be the case that S sees that *p*, *or* that she is in some deceptive state that she is not in a position to introspectively discriminate from a state of seeing that *p*; either case can be correctly described as one in which it perceptually seems to S that *p*. I take it that this much is not controversial, since it just is a theoretically neutral description of seemings. What makes disjunctivism a substantial and hotly debated thesis is the claim that the two disjuncts do not pick out the same fundamental experiential kind: there is a fundamental difference between mental episodes that are episodes of seeing and mental episodes that are merely indiscriminable from seeings.

Disjunctivism says that there is a fundamental divide between veridical and falsidical experiences despite their subjective, qualitative indiscriminability. This makes it a very

²⁵ I do not mean to say that there are two distinct, let alone incompatible, versions of NR. The point is just that different naïve realist authors emphasize different aspects of the view.

²⁶ Fish (2009) adopts a radical solution, for he thinks that hallucination merely *seems* to have a phenomenal character, while in fact it has none.

unorthodox view when compared to a traditional, broadly Cartesian, picture of the mind. Let us consider again the argument from hallucination:

Argument from Hallucination

P1: In cases of hallucination it appears to one that one is immediately aware of a mind-independent object *o* that has a sensible quality *F*, although in fact there is no object at all;

P2: If it appears to one that one is immediately aware of a mind-independent object that has a sensible quality *F*, then there is an object that is *F* of which one is immediately aware (Phenomenal Principle);

P3: In cases of hallucination there is no mind-independent object that one could possibly be immediately aware of;

P4: Either one is immediately aware of a mind-independent object or one is immediately aware of a mind-dependent object;

C1: In hallucination one is immediately aware of a mind-dependent object (sense datum/data);

P5: Cases of hallucination can be subjectively indiscriminable from cases of veridical perception;

P6: All subjectively indiscriminable experiences are experiences of the same fundamental kind (Common Kind Assumption);

P7: Perceptual experiences of the same fundamental kind have the same kind of immediate objects;

C2: Both in perception and hallucination one is immediately aware of mind-dependent objects (sense datum/data).

Subjective indiscriminability is taken to be of consequence in this argument, for P6 says that indiscriminable experiences are experiences of the same fundamental kind, and it is only because of that that we get to C2 in the end. Notice how at least some representationalists would agree with this step of the argument (although they block it at an earlier stage by denying P2), thereby attributing a certain importance to indiscriminability. Disjunctivists see the matter differently. Although they would typically deny P2, just as all representationalists do, this is not the hallmark of disjunctivism. The essential disjunctivist point is rather the denial of the *Common Kind Assumption* (CKA) expressed by P6, namely the idea that

subjective indiscriminability is of consequence for the metaphysics, the phenomenology, or the epistemology of perceptual experience in general:

Philosophy has tended to credit a certain kind of significance to the fact that a pair of states can both be truly described as states of its looking to S as if things are some way (where the way is the same in both cases, and the looking is phenomenological). One such significance is that S' experientially acquired reasons for belief must be the same in both cases. Another is that the intrinsic nature of such states must be the same in both cases, and cannot involve anything "outer". And another is that, in both cases, experiential phenomenology, and what explains it, must be the same. *It is a mark of disjunctivism about experience to question these assumptions.* (Haddock & Macpherson 2008: 21, emphasis mine)

The disjunction in terms of which disjunctivists conceive of the dialectic between veridical and falsidical experience can be articulated in at least three (*non-mutually exclusive*) ways, depending on the kind of significance that indiscriminability is denied to have. An influential form of disjunctivism holds that the disjunction in question adverts to a metaphysical difference, in that veridical and falsidical experiences do not belong to the same fundamental kind; call this *metaphysical disjunctivism*. If metaphysical disjunctivism is correct, the fact that either of a pair of experiences can be correctly described as it looking to S that *p* does not imply that the two experiences are instances of fundamentally the same kind. Another version of disjunctivism holds that veridical cases have a different epistemological import, that is, they don't give us the same grounds as falsidical cases for acquiring perceptual beliefs; instead, they provide factive reasons — let us call this *epistemological disjunctivism*. Finally, yet another version holds that veridical cases have a different phenomenal nature; call it *phenomenal disjunctivism*.^{27 28}

²⁷ For metaphysical disjunctivism, see Brewer (2008, 2011) and Martin (2004). For epistemological disjunctivism, Pritchard (2012) and McDowell (1982). For phenomenal disjunctivism, Martin (2002) and Brewer (2008, 2011). For two collections on disjunctivism in general, see Byrne and Logue (2008), and Haddock and Macpherson (2008).

²⁸ Disjunctivism could also be formulated in terms of object-perception, rather than seeing that *p*: thus the relevant contrast would be one between *either* seeing an object *o*, *or* being in some state subjectively indiscriminable from seeing *o*. Depending on the kind of disjunctivism that one wants to uphold, one would then have to point out that the two disjuncts have a different metaphysical, or phenomenal, or epistemological significance. At least some disjunctivists may be more comfortable with a formulation in terms of object-perception: for example, Brewer thinks we see objects, not facts; furthermore, *seeing that* is often used to denote a *post-perceptual* epistemic state, although there are exceptions (see McDowell 1994 (26) and Fish 2009 (51—54)).

What exactly is the relationship between NR and disjunctivism? A plausible way of construing their relation is that disjunctivism is a dialectical and theoretical resource that one can advert to in order to uphold a naïve realist theory of veridical perception. Thus Soteriou buys into this way of setting up the dialectic:

Disjunctivists are often naïve realists, who hold that when one perceives the world, the mind-independent objects of perception, such as tables and trees, are constituents of one's experience. In other cases, such as hallucinations, it seems out of question that such objects are constituents of one's experience. It follows that on a naïve realist view, the veridical perceptions and hallucinations in question have a different nature: the former have mind-independent objects as constituents, and the latter do not. (Soteriou 2009: 1)

Soteriou suggests that analysing in terms of a disjunction a pair of indiscriminable experiences (one of them being hallucinatory) is the natural consequence of a naïve realist perspective, which demands indeed that hallucinations have a different nature from veridical perception. In other words, NR is the *motivation* for being a disjunctivist in the first place.^{29 30}

Earlier in this section, I labelled sense-data theories and to a certain extent representationalism scientific-like approaches, meaning that they start with the theoretical problem of illusion and hallucination, and then revise the common-sense conception of perceptual experience in light of the response that they give to that problem. NR/disjunctivism, by contrast, holds that cases of perceptual error should be accounted for derivatively with respect to perception proper. The way in which the NR/disjunctivism pair defines the form of the inquiry is the reverse of the representationalist and sense-data theories', for the question that naïve realists/disjunctivists ask is "how, given that perception *is* openness to the world, should we account for illusion and hallucination?". This makes NR/disjunctivism a piece of therapeutic philosophy, in the sense that it aims at removing supposedly theoretically-contrived obstacles to what its proponents think is the most natural, intuitive and introspectively evident description of perception.

The therapy recommended by NR/disjunctivism might seem cheap, though – nothing more than a dialectical ploy: veridical perception falls on the good side of the disjunction,

²⁹ Martin (2002: 396-396) and Logue (2013: 107) make the same point.

³⁰ I leave it open whether all kinds of disjunctivism are motivated by NR: while this seems certainly the case for metaphysical and phenomenal disjunctivism, epistemological disjunctivism may be motivated by different considerations (see for example Pritchard 2012).

thereby being an instance of the fundamental kind of perceptual experience that NR takes as its explanatory target, whereas misleading experiences fall under the bad disjunct, and as such are instances of a fundamentally different kind of experience. Whether a greater degree of substantiality can be achieved through this explanatory scheme depends on what NR/disjunctivism can actually say on illusion and hallucination. In particular, what I want to investigate is whether a brand of NR that dispenses with content can successfully account for illusion. In the next section, I briefly comment on hallucination and explore in detail the question of illusion, examining several purely relationalist proposals (that is, versions of NR that dispense with content) on the market and argue that none of them is entirely satisfactory, although Brewer's account is partly successful.

4. Pure Relationalism and Perceptual Error

According to the characterization that I have offered, when pure relationalists talk about perception as openness to the world they have in mind several interrelated themes: one such theme is that perceptual experience does not act as a source for the selection of one fact from among many as to how any one given object looks – that was the topic of chapter two. Another theme, that has emerged over the course of this chapter, is that perception is a passive confrontation with reality and, as such, is innocent with respect to representational error. Yet another theme is that perception is to be phenomenologically and metaphysically characterized with reference to the objects that one is aware of in perceptual experience.

How do pure relationalists explain, in the light of these fundamental tenets of theirs, the possibility of misleading experience? Naturally, it is important to notice that illusion and hallucination may be treated separately: in the case of illusion one may still want to claim that the experience in question is fundamentally a perceptual relation between a subject and a material object, while this is obviously not possible in the case of hallucination. In what follows I briefly comment on the case of hallucination and then examine illusion in far greater detail.

The most popular and at the same time controversial account of hallucination from a naïve realist standpoint has been formulated and defended by Mike Martin, and versions of it have been endorsed by Bill Brewer and William Fish. The account in question is an *epistemic* conception of hallucination, which holds that hallucination is to be characterized purely in

negative terms, as an experiential condition which is not introspectively discriminable from a corresponding veridical perception. According to Martin, all that there is to having an hallucinatory experience to the effect that there is a red ball in front of one is to be in a state which is subjectively (that is, introspectively) indiscriminable from one that is a state of seeing a red ball in front of one; thus hallucinations are states “whose only positive mental characteristics are negative epistemological ones – that they cannot be told apart by the subject from veridical perception” (Martin 2004: 73-74).

There is a reason why Martin thinks the naïve realist should only give a negative characterization of hallucination, that is, one that only characterizes hallucination with reference to veridical perception, rather than in terms of positive mental characteristics. For suppose that naïve realists were to identify a positive feature of hallucination; for example, they might characterize a hallucination of a yellow banana as a state whereby a subject perceptually represents that there is a yellow banana. Plausibly, then, the behavioural, phenomenological and epistemological aspects of that hallucinatory experience would be explained in terms of the previously identified positive feature, namely a representational content. Also plausibly, there is a neural state with a proximal neural cause that is sufficient to cause the subject to go into the state in question. That is, the subject being in the relevant neural state would be sufficient for her to have an experience with certain psychological features, which would be explained in terms of a representational content. Now, it certainly is legitimate to suppose that the very same proximal cause that is sufficient for the subject to hallucinate a yellow banana could be the last (that is, the proximal) link in a causal chain that originates from an actual yellow banana and triggers a veridical visual experience of a yellow banana. What would explain all the phenomenological, epistemological and behavioural features of the subject’s experience in the veridical case? By naïve realist lights, the fundamental *explanans* would be the fact that the subject is in a perceptual relation to a yellow banana. However, the very same psychological features could have been brought about solely in virtue of the proximal neural cause, had the subject been hallucinating. And that very same proximal cause is also present in the case of a veridical perception of a yellow banana; therefore, it seems that the fundamental explanatory role is being played by the proximal cause and the representational content associated with it, *both* in the case of veridical perception and in the case of hallucination. As a result of Martin’s work, this point is often expressed by saying that whatever positive feature of hallucination were invoked to

explain some relevant psychological features of hallucinatory experiences, would *screen off* the perceptual relation to the environment from being the fundamental *explanans* in the case of veridical perception. Matthew Nudds summarises this clearly:

The following screening off principle seems plausible: an event *e* being *h* screens off being *p* as an explanation of some effect of *e* if, had *e* occurred and been *h* but not *p*, then *e* would still have had that effect. For example, a cloth sample being red screens off its being scarlet as an explanation of its being picked out by a machine if, had the sample been red but not scarlet (by being crimson, say) it would still have been picked out. In that case, it would be false to say that it was picked out because it was scarlet since it would have been picked out whether or not it was scarlet. (Nudds 2009: 338, footnote omitted)

Suppose that an event *e* is a hallucination *h* of a yellow banana, and that we explain its psychological effects in terms of a representational content. Suppose that another event is a veridical perception *p* of a yellow banana. Given that the proximal cause which is sufficient for a state to be *h* would also be present in *p*, it follows that an event can be both *h* and *p*. But if the basic psychological features or effects of *e* would have been present even if *e* had not been *p*, then that effectively screens off *p* from being the most fundamental *explanans* of the psychological features and effects of *e*. To draw the moral, it seems that NR is in trouble as soon as its proponents identify a positive feature of hallucination.

The trouble does not arise if one gives a purely negative characterization of hallucination, though. Suppose that the only positive property of a hallucination of a yellow banana is the negative epistemic property of being subjectively indiscriminable from a perception of a yellow banana. This, it seems, is sufficient to explain all the psychological effects that the hallucination has on the subject (it seeming to the subject that there is a yellow banana, the subject's forming a belief to the effect that there is a banana, and the subject's trying to reach out for a banana). Crucially, unlike the individuation of a positive feature of hallucination, the negative property of being indiscriminable from a veridical perception does not screen off the perceptual relation from being the fundamental *explanans* of all the psychological effects of the experience on the subject in the veridical case. The reason why is that there is an asymmetric dependence of the effects of hallucination on the effects of veridical perception:

Why did James shriek like that? He was in a situation indiscriminable from the veridical perception of a spider. Given James's fear of spiders, when confronted with one he is liable so to react; and with no detectable difference between this situation and such a perception, it must seem to him as if a spider is there, so he reacts in the same way. (Martin 2004: 68)

James shrieked because he was in a situation that was subjectively indiscriminable from a veridical perception of a spider. There are two ways in which he could have been in that situation: by actually being in a situation which was a perception of a spider, or by being in one that was a hallucination. Now, if the property of being indiscriminable from a veridical perception of a spider screened off a veridical perception of a spider from explaining why James shrieked, then the property of being indiscriminable from a veridical perception should also lose explanatory power: for the property of being indiscriminable from a veridical perception could explain certain psychological effects only under the assumption that *being* a veridical perception was something that could have explanatory force in the first place; to put it briefly, James being in a state that is subjectively indiscriminable from that of seeing a spider could not explain why he shrieked if his actually seeing a spider did not have a prior explanatory power. As a result of Martin's argument, it is often concluded that a naïve realist should answer the question of how to explain hallucination by invoking negative disjunctivism.

In this thesis one will find little material on hallucination, and this may require an explanation. The reason for that is actually quite simple: this is a work on perception as openness to the world, and while illusion may be classified as an instance of openness to the world insofar as there is a mind-independent object that one sees, this does not apply to hallucination. I have no particular qualms about giving a disjunctive treatment of the perception/hallucination pair (although I do raise one minor worry about negative disjunctivism in 4.3.2), and this does not require major adjustments for the view I want to articulate. In effect, I only need to emphasize a consequence which was already implicit in my endorsement of singular contents. The consequence in question is that introspective indiscriminability turns out to be less important than it might initially supposed to be when starting from a sense-data theory or from certain versions of representationalism. Indeed, if the contents of perception are singular, indiscriminability does not entail sameness of content the way, because two experiences of two qualitatively identical red spheres might be

introspectively indiscriminable, although one makes singular reference to *this* ball, and the other to *that* ball. The case of hallucination simply brings into relief the same point from a different angle. A hallucination of a red sphere might be introspectively indiscriminable from a perception of a red sphere, but that does not imply that there are no fundamental differences between perceptions and hallucinations *qua* experiences. Whether the relevant difference should be captured by saying that the content of perception is singular while the content of hallucination is existential, or by saying that hallucinatory experiences make singular reference to non-existent objects (Smith 2002: 237), or that in veridical perception we represent objects in virtue of seeing them whereas in hallucination representation is bedrock (Logue 2013), or by endorsing the more radical negative disjunctivist account offered by Martin (2004, 2006), is not something I am overly preoccupied with in this thesis – but if I had to pick a side, I would say that hallucination is an intentional state directed at a non-existent object, which is in line with the general conception of intentionality that I have offered in chapter one.

In what follows, I examine three possible approaches to illusion, with close attention to the sophisticated account developed by Brewer (2011), and I argue that none of them is entirely satisfactory, although I acknowledge that Brewer's theory is partly successful.

It will be convenient to phrase the challenge for purely relationalist views in terms of the subjective, phenomenological aspects of experience. Illusion presents a difficulty for the thesis that the phenomenal character of perceptual experience is constituted by the object that the experience is of, because paradigmatic cases of illusion are precisely situations in which things appear to the subject to be different than they are. If, for example, it seems to one as if there is a red ball in the room although the ball is in fact blue, it is certainly not a straightforward matter to articulate a sense in which the experience is metaphysically and phenomenologically constituted by the ball and its blueness. In the face of this problem, there are three options that I want to consider for pure relationalists: one is to introduce some extra-relational, representational or qualitative factor to account for the illusory aspects of the experience (*mixed accounts*). Another is to extend the perceptual relation so as to encompass *appearance properties*, and use these to explain illusion. Yet another (Brewer's account) is to make do with the object and its intrinsic properties. I consider these options in turn, and argue that the third is by far the most promising, although not completely exempt from phenomenological and metaphysical problems, which I discuss toward the end of this chapter.

Finally, I should remark that Brewer's account also raises problems related to the epistemic and phenomenal determinacy of perception, that will be subject to thorough examination in the fourth chapter, in connection with the topic of particularity and generality in perception.

4.1 *Mixed Accounts*

A mixed account of illusion is one on which an illusory experience is partly constituted by the object that the experience is of, and partly by some other factor, such as a representational content to the effect that the object has a certain feature which in fact has not. Mixed accounts have been very convincingly criticised by Smith (2010), and I think that the prospects for such views are by no means good, due to metaphysical and phenomenological problems. Let us suppose, following Smith (2010: 388-390), that one sees a green square but misperceives it as yellow. A mixed view would hold that the very same square that the subject sees accounts for the veridical aspects of her experience, and some other factor, such as a representational content or *quale*, accounts for the yellow illusory appearance.

This faces insurmountable problems, though. For one does perceive veridically the square shape of the object precisely in virtue of seeing the distribution of the colour that one misperceives. It is indeed a familiar point, much insisted upon by Husserl (2001), that the shape, extent and colour of an object are inextricably linked, both as features of the object and as *perceived* features, in the sense that one cannot perceive the extent or shape of an object without perceiving its colour (and vice versa). But how could a mixed account of illusion explain this fact? On the view under examination, the veridical perception of the square shape is metaphysically and phenomenologically constituted by the very square one sees. However, one could not veridically perceive the square shape in the first place, if it were not for the illusory experience of the colour, which, by hypothesis, is not constituted by the actual colour of the object. Therefore, the mixed account theorist has to embrace a very unstable combination of claims, for she has to say that (i) the square shape partly constitutes one's experience, although (ii) that in virtue of which the shape is perceived (namely the apparent distribution of the colour) does not.

This combination seems untenable, because whatever notion is invoked to explain the phenomenal character of one's experience of the colour of the object, must also be invoked in explaining the phenomenal character of one's experience of the square shape, precisely

because perceiving the shape of an object is a matter of perceiving the distribution of its colour(s). The interrelatedness of colour, shape and extent also explains why it would not work to say that veridically perceiving the square is a matter of correctly perceiving the *distribution* of the colour, *but not the specific colour* that the surface has. This move would be desperate from both a metaphysical and a phenomenological point of view. Metaphysically speaking, the colour of a surface can be merely abstracted from the surface that it qualifies, but not actually separated in the sense in which one of the legs of a table can be detached from it – this is the point of the Husserlian distinction between *moments* and *parts* of an object.³¹ *A fortiori*, phenomenologically speaking the shape of a surface is perceived by perceiving the particular colour that the object has, and not just *some* colour or other. The proponent of the mixed account, then, is in a predicament, because if the experience of the apparent yellowness of the square is not constituted by the actual colour of the square, so is not the experience of its shape either – which is in plain contradiction with the basic insight that NR starts with. As Smith points out, the problem generalises:

...whatever bad factor is postulated to account for an illusorily appearing feature in a partially illusory perceptual experience will undermine the naïve realist account of any veridically appearing aspect whose appearing presupposes and is fully determined by the appearing of such a feature (Smith 2010: 389).

The issue becomes even more evident if one supposes that both the shape and colour of an object are misperceived. For example, we might suppose that we misperceive a green square as yellow and rectangular. By naïve realist lights, neither the shape nor the colour of the object could constitute the subjective character of one's experience. But then it becomes hard to spell out any intelligible sense in which the object in question could be a phenomenological and metaphysical constituent of one's experience of it.

It seems that a mixed account of illusion is a non-starter, because it results in an unstable hybrid that cannot even respect the fundamental naïve realist insight that experiences are phenomenologically and metaphysically constituted by the objects that they are of.

³¹ See Husserl 2001, 3rd investigation.

4.2 Appearance Properties

Louise Antony (2011) has put forward a view to the effect that in cases of illusion we do see actual features of objects or scenes, only of a kind different from the intrinsic properties that we are committed to attributing to objects through our ordinary discourse and practices. What we see in cases of illusions is, according to Antony, a kind of relational features, often called *appearance properties*. Of a straight stick that looks bent because partly immersed in water, for example, Antony says:

What I am denying is that the feature of our experience that we try to capture with the locution “It seems like . . . is bent” is a feature that represents the objective shape-property of being bent. The feature that is actually present in our experience of a straight stick in water is a feature that infallibly corresponds to a property the stick actually does possess, namely the objective feature of objective sticks that triggers visual systems like ours to produce the particular kind of percept it produces (...call it *bentishness*). Our experience, therefore, does not have the content “the stick is bent.” It rather presents to us a bentish stick, something it cannot fail to do, given the openness of perception and the manifest bentishness of the stick. (Antony 2011: 34-35)

Whatever one might think of the metaphysical status of appearance properties, it is not clear that this proposal can keep faith to the conception of perception as openness to the world that is put forward as the main motivation for it. For even if one concedes that a straight stick becomes bentish when partly immersed in water, it certainly is the case that it also remains straight, since water does not affect the shape of a stick. But then there are two properties involving the shape of the stick, both of which may be relevant to perceptual phenomenology: *straightness* and *bentishness*. It seems that Antony should say that only the latter figures in the phenomenology of one’s experience. However, if perception is mere confrontation with the actual constituents of reality, and if the phenomenology of perception is a question of what one is sensorily aware of, then it seems she should say that both *straightness* and *bentishness* figure in one’s experience.

Antony cannot say that a certain property (*bentishness*) is selected over another (*straightness*) as that which figures in the experience, for that would be a mere notational variant of the representational strategy that consists in saying that the stick is represented as

bent. So, if Antony wants to hold that only *bentishness* determines the phenomenology of the experience, she has to find some consideration that speaks against one seeing *straightness*.

On a plausible conception of what it is to see something (an object, or a feature of an object), there are two conditions that need to be satisfied in order for claims of the form “S sees *o*” or “S sees *o*’s F-ness” to be true: the experience that S has must be causally dependent on the feature or object that is being seen (Grice 1961, Strawson 1979), and the experience must be such that it discriminates the object or feature in question from its background (Dretske 1969). However, it doesn’t seem that there are good grounds for denying that either of these two conditions obtain in the case of the stick in water.

Suppose one said that the relevant experience depends on *bentishness*, not on *straightness*, thus denying the satisfaction of the first condition. This would be clearly wrong, for the stick in water is bentish, and looks bent, *because* it is straight (and partly submerged in water, of course). Therefore, the causal role played by the intrinsic property of the stick in triggering the experience is, if anything, more fundamental than that played by the relational property.

Could Antony say that in the relevant experience *straightness* is not discriminated from the background, and hence that the second condition on seeing is not satisfied? That doesn’t seem right either. For cases where we fail to see something because the experience doesn’t discriminate it are typically cases in which the object looks too much like the background. Thus, we may fail to make out a white sheet of paper stuck on a white wall, or recognize a chameleon that has taken on the same colour as the foliage surrounding it. In such cases, we may want to say that, although the objects in question do stand in the appropriate causal relation to us insofar as they reflect light onto our retinas, they are not seen, because we are not in a position to discriminate them from their background. However, there is nothing analogous going on in the case of the stick in water. Here the background is an expanse of water, and it certainly does not prevent us from discriminating the stick or some of its features.

In addition, it is not clear that Antony’s proposal can make sense of illusory experience being somehow a less than optimal kind of perception – that is, one in which we or our visual

system get something wrong about the way things are.³² For according to Antony, we do perceive a property that the object really has, so there seems to be no error at all. Increasingly, it also seems to us as if we perceive the *intrinsic* bent shape of the stick, and not some relational, appearance-property that is somehow related to its intrinsic shape. Given the purely relationalist standpoint that Antony adopts, this outcome is paradoxical, for it exposes her account as a kind of error-theory: it seems to us as if we perceive intrinsic properties, but in fact we perceive a whole different *kind* of items, namely appearance properties; it hardly needs be pointed out how close this comes to the first step in the argument from illusion, whereby it is claimed that although it seems to us as if we perceive directly mind-independent objects, the entities that we are primarily aware of are mind-dependent sense-data.

It seems that Antony's account of illusion, apart from involving a commitment to the metaphysical thesis that there are appearance-properties that is not argued for, is affected by at least three problems: first, it seems unable to provide a justification (other than in representational terms) for the claim that we see appearance-properties rather than intrinsic properties in cases of illusion. Second, Antony cannot make it intelligible why illusions are less than optimal cases of perception. Third, she is committed to a kind of error-theory, which obviously does violence to her naïve realist inclinations.

4.3.1 *Visual Similarities*

So far, I have deliberately ignored what I consider in fact to be the best chance for pure relationalists, namely to stick to their guns and claim that, even in illusory experience, the way things are for the subject is to be explained purely in terms of that subject being in

³² Logue (2009) makes the bold suggestion that, according to (what I call) pure relationalism, in an illusory perceptual experience of an object, "the illuded subject doesn't misperceive the object at all" (70), because she perceives a "property that the object *really has*" (70). If Logue is right, my objection to pure relationalism simply begs the question, as the relationalist point is precisely that illusion is not a sub-optimal perception.

Although I find this suggestion very counterintuitive, I admit that it is not obvious how to refute it: if we suppose that Antony could have a good explanation as to why in the illusory case the experiencing subject only sees *bentishness*, not *straightness*, we get to a picture on which she does not misperceive anything, because not seeing a property that an object has does not count as misperceiving it. According to this line of reasoning, that illusory experience is not sub-optimal is simply the outcome of a correct analysis, not a problem for pure relationalism.

However, even setting aside my worries about the thesis that the subject does not see *straightness* in the relevant case, there clearly is a respect in which illusory experience is sub-optimal: illusion is *epistemically* misleading, for if one were to judge the shape of the stick based only on experience, one would judge it to be bent (not bentish!); the simplest explanation for this is that the falsity of the judgement reflects the fact that there *is* something less than optimal in the experience itself — in other words, the best explanation for the subject's perceptual judgement that the stick is bent is that the experience perceptually represents the stick as bent.

relation to an object and its intrinsic features under specific circumstances. This is an option that tends to be implicitly ruled out right at the outset, for it is taken for granted that when one sees, say, a green square that appears yellow, the subjective yellow appearance “*clearly*, must be accounted for by something other than the green square being a constituent of the experience” (Smith 2010: 388, emphasis mine). But perhaps we should not accept this assumption prior to scrutiny; would it not be possible to spell out a sense in which the actual greenness of the square is a constituent of the experience, even if the square appears yellow?

There is a somewhat crude version of this relationalist line that does seem ineffective. Travis writes that one’s perceptual phenomenology “lies in what one witnesses” (Travis 2013: 840). So whatever it is that one sees, that is what explains the phenomenology of one’s experience, even when the way things are and the way things look to a subject come radically apart – for example, what Pia sees, and hence constitutes the phenomenology of her experience, might be John walking, although for some reason all looks like “a fleeting blur to her” (Travis 2013: 839). However, if this – “what one witnesses” – is all what the pure relationalist has to go on, then it is no longer clear that it is a theory of perceptual *experience* that is being put forward, for one is left wondering how merely gesturing toward the object or event that the subject is in relation to is supposed to account for the phenomenological dimension of perception. Susanna Siegel has made a similar point in terms of a distinction between coarse-grained and fine-grained awareness of properties:

On a fine-grained interpretation, when you are aware of John walking, the property of walking characterizes how John looks to you. On a coarse-grained interpretation, you can be aware of John walking, without the proposition that someone is walking being so much as suggested by your experience, as something reasonable to believe on its basis. Sid and Pia might both be perceptually aware of exactly the same visible instances in the coarse-grained sense, yet have different fine-grained awareness of those instances. This difference is a difference in whether the properties whose instances you’re aware of characterize how the things you’re aware of look (that is the fine-grained option), or whether they simply identify which things these are (that is the coarse-grained option) (Siegel 2013: 859).

Siegel’s distinction between fine-grained and coarse-grained awareness brings out nicely the difficulty that purely relationalist theories face. If Pia sees John walking, but all looks just like a fleeting blur to her, then the awareness that Pia has of John walking is coarse-grained, in the

sense that the individual *John* and the property *walking* merely identify what is seen, but fail to characterize the way things are with the subject of the experience. The challenge for the pure relationalist is that of articulating a sense of constitution on which the constituents of the experience characterize how things appear to the subject – in other words, square constitution with a fine-grained awareness of objects and properties.

If any purely relationalist theory of illusion has a chance of meeting this challenge, it must be the rich and sophisticated account offered by Bill Brewer (2011). Here is the core of it:

...in a case of visual illusion in which a mind-independent physical object, *o*, looks F, although *o* is not actually F, *o* is the direct object of visual perception from a spatiotemporal point of view and in circumstances of perception relative to which *o* has visually relevant similarities with paradigm exemplars of F although it is not itself actually an instance of F. (Brewer 2011: 105)

The main idea is quite simple: objects may look different than they are in virtue of visual similarities to other objects, similarities which they have or come to have under particular circumstances. The conscious registration of these similarities makes it intelligible how an object may appear to be a certain way to a subject, even if in fact it is not. More in detail, there are three key notions that constitute the account: the notion of visually relevant similarities, that of paradigm exemplars, and the role of the spatiotemporal point of view along with the circumstances of perception.

An object, it is said, might look F even if in fact is not an F, due to its visual similarities to objects that actually are Fs. This immediately raises the question of how we should understand visually relevant similarities, given that pretty much anything can be seen as having some similarity with this or that object; how are we to determine which similarities are relevant? Brewer's response consists in articulating the notion in empirical terms: if any two objects or configurations of reality share sufficiently many properties among those which play a significant role in visual processing, then they do have visually relevant similarities; that is, as Brewer says, "visually relevant similarities are similarities by the lights of visual processing of various kinds" (Brewer 2011: 103). For an object to look F, it must also be the case that the F-object that the object of the actual illusory experience putatively looks like be a *paradigmatic exemplar* of an F-object, where paradigmatic F-objects are those exemplars of

Fs that play a special role in the acquisition of the concepts associated with the expressions that we use to refer to the F-objects – particular such objects as, say, a red ball or a red mail box may be.

Finally, a crucial role is played by the spatiotemporal point of view and the circumstances of perception (lighting conditions, but also subjective factors like attention, fatigue and short-sightedness), since an object may have visually relevant similarities to an F-exemplar only because it is seen under specific circumstances – perhaps under water, or through distorting lenses, or under a particular illumination, and so on. The spatiotemporal point of view, along with other circumstances of perception, is an integral part of what it means to be acquainted with an object, for we always perceive objects in particular sense modalities, from a particular point of view and under particular circumstances. Perceptual consciousness is thus characterized as a three-place relation that comprises the subject, the object, and the spatiotemporal point of view along with the circumstances of perception.

Brewer's account can be best appreciated in its direct application to concrete examples of illusions. Let us consider again the ML diagram. As we've already seen, the two main lines look to be different in length, while in fact they are exactly the same length, and the impression of inequality is due to the inward and outward hashes. According to Brewer, the ML diagram has visually relevant similarities "with a pair of lines, one longer and more distant than the plane of the diagram, one shorter and less distant" where this latter pair of lines would be "a paradigm of inequality in length" (Brewer 2011: 102). Therefore, the idea goes, it certainly is intelligible how, upon seeing the ML diagram, one may register those visual similarities and "take that very diagram as consisting of unequal lines" (102). And yet the particular conscious condition that it is to see an ML diagram still is elucidated as fundamentally a relation of acquaintance with that very diagram, without any need to appeal to sense-data, perceptual representations, or special appearance-properties over and above those properties that we are committed to attributing anyway through our ordinary discourse to the things we see.

Likewise for illusions involving shape. A straight stick in water is visually similar to a paradigmatic bent stick that has "its top half coincident with the unsubmerged half of the seen stick, and its bottom half in the position of the relevant virtual image of the bottom half of the latter from the subject's point of view and given the refractive index of the liquid in question" (Brewer 2010: 106). The sense in which the actual stick and the hypothetical

paradigmatic bent stick are similar is eminently visual, for the idea is that “in the region of space in the vicinity of the eye...light from corresponding parts of the two sticks travels, or would travel, along the same paths” (106). As in the case of the ML diagram, it is a fundamental relation of acquaintance with the object just as it is, which makes it intelligible how one may register the similarities that the stick has to a bent stick and hence take it to be one; as in the ML diagram case, no reference is made to sense-data, representational contents or appearance-properties.

Finally, the account seems also adequate for handling illusions involving colour. A piece of white chalk under a red light is visually similar to a paradigmatic piece of white chalk. As in the case of the straight stick in water, it is the spatiotemporal point of view and the circumstances of perception which play a crucial role in elucidating the sense in which the two pieces of chalk are similar, for it is in virtue of the particular lighting conditions that the two pieces of chalk come to have visually relevant similarities, which in this particular case consist “in the similarity of the light reflected from both” (106). Thus everything is explained purely in terms of a three-place relation of acquaintance that comprises the experiencing subject, the object just as it, and the spatiotemporal viewpoint along with the circumstances of perception.

Brewer’s theory seems to have at least two important virtues. First, it seems the only purely relationalist view that could square naïve realist claims about metaphysical and phenomenological constituency with the obvious datum that the things we see sometimes look different than they are. The exchange between Siegel and Travis that we saw earlier in this subsection issued in a challenge for the pure relationalist to elucidate a sense in which perception is to be characterized solely in terms of what one is sensorily related to, that could also be a sense in which those things that one is related to characterize in a fine-grained way one’s awareness of them, thus not merely identifying what one sees. In this connection, there is no doubt that Brewer takes up the challenge. Second, the theory appears to be suitably flexible, in the sense that it can be applied equally well to illusions involving size, shape and colour.

Brewer’s account of illusion has been placed under close scrutiny by Smith (2010), who claims that there are several kinds of illusions for which Brewer’s view is inadequate, because it is incapable of showing how the direct objects of illusion together with their sensible qualities can figure in one’s experience in a phenomenologically intelligible way. In the next

subsection, I offer a discussion of Smith's objections together with possible responses on Brewer's part. I argue that some of the objections that Smith raises might be misguided, and thus I partially vindicate Brewer's relationalist account. However, I also argue, partly relying on Smith (2010), that there is a problem for Brewer concerning after images, and that that, in turn, may cast a shadow on the negative disjunctivist account of hallucination. But before I start, I should emphasise that I consider Brewer's account of looks and illusion one of the best currently on the market, and most of my critical remarks in this chapter are to be understood more in the way of a general discussion of problems that are *anyone's* problems, rather than in the way of a criticism of Brewer's account. My real objection against Brewer comes in the fourth chapter: there I argue that the only way Brewer's account can work, insofar as it does work, is by collapsing into a version of representationalism.

4.3.2 *Phenomenological Intelligibility and Sensible Qualities*

Let us draw, following a standard practice in psychology, a distinction between physical and physiological illusions on the one hand, and cognitive illusions on the other.³³ Physical illusions are, roughly, those that depend entirely on the laws of optics and the way these determine the retinal images impressed on our retinas. Thus, given the refractive properties of water, it is purely a physically-determined fact that a straight stick partly immersed in water should impinge on our sensory system so as to look bent.

Physiological illusions are those whose illusory character is not necessitated by particular configurations of reality or the laws of projective geometry, but rather require an explanation that takes into account the internal working of our visual system. A popular example involves a picture of the American flag with reversed colours (yellow, black and green), with a white spot in the centre. If you stare at the white spot at the centre of the flag for about thirty seconds and then look at a white surface, you will have an experience as of a projection of the flag in its original red, white and blue colours. This depends of course on the particular way in which colours interact with the physiology of our visual system, and thus cannot be explained purely in terms of the features that the picture has.

Finally, cognitive illusions are those that cannot be explained purely in terms of what hits the eye together with the physiological peculiarities of our sensory apparatus. The ML

³³ See for example Fish (2009) and Gregory (1997).

diagram is a case in point: given the two lines, their actual length and the additional hashes, there does not seem to be enough in the laws of optics or in the physiology of our visual system to explain why they look different in length. Therefore, there must either be some assumption that our visual system makes in virtue of which it categorizes the diagram as a case of inequality in length, or some unconscious psychological association that we as subjects make between ML and cases of inequality in length. Likewise, a coil of rope may look like a snake to the ophidiophobic subject, although nothing in the way things are or in the internal working of her visual system necessitates this particular experience. Therefore, some psychological association between the visual features of the scene and snakes must be brought to bear on the explanation of the illusion.

Smith (2010) argues that Brewer's approach might be effective with cases of cognitive illusions, but on close inspection faces serious problems in the case of physically or physiologically-based illusions. First of all, it should be noticed that the role of the object in shaping the experience is much less straightforward in cases of physical illusion than it is in cognitive and physiological illusions. With a cognitive illusion such as ML one might hold, as Brewer does, that the experience is shaped by the actual diagram, although one might intelligibly take that very same diagram to be an instance of inequality in length, due to a psychological association with a pair of unequal lines at different depths. Brewer is indeed in a position to offer some elucidation of the sense in which we see the two lines exactly as they are. He points to two facts: first, if the misleading hashes were to gradually shrink in size and then vanish, one would not experience the two main lines as gradually changing in length. Rather, one would "gradually come to realize that any previous inclination to take them to be unequal in length was mistaken, as the power of the hashes to mislead in this way diminishes" (Brewer 2006: 170). And this suggests that one had been "presented in experience throughout with the very same pair of lines, equally extended in space as they actually are, whose unchanging identity in length becomes gradually more apparent...as their similarities with an alternative configurations of unequal lines at different distances become less salient" (Brewer 2006: 170-171). Second, Brewer notices that the endpoints of the main lines seem to be exactly where they are, even when the misleading hashes are present. Therefore, the suggestions goes, it is perfectly sensible to say that the particular conscious condition that consists in seeing the ML diagram should be elucidated simply by reference to that very same

diagram, without the need for positing dubious entities such as sense-data or theoretical notions such as content.

Things become murkier when we consider physiological illusions, though. Consider a case of visual after-effect such as that of the picture of the American flag, or a simpler one in which a white surface looks red because one has been staring at a green surface beforehand. It seems that a white surface and a red surface are not similar in the same sense in which the ML diagram is similar to a pair of unequal lines. In effect, it just seems that they are not similar period; white and red are two very different colours. Brewer gestures toward a possible account of afterimages by saying that they are a case of a hallucination which is superimposed on a mind-independent object of acquaintance. Thus when one sees a blue afterimage on a white screen after staring at a bright red stimulus, this should be explained in terms of a presentation of the screen as a mind-independent object, “supplemented by a hallucination introspectively indistinguishable from seeing a blue stimulus corresponding to the original one” (Brewer 2011: 115).

Brewer’s account has the interesting consequence that afterimages are to be treated as cases of seeing, for the only positive feature of hallucination is, according to negative disjunctivism, that of being introspectively indiscriminable from an episode of seeing. This raises some perplexities though, as afterimages typically have a distinctive phenomenology, which cannot be adequately captured in terms of the phenomenology of *seeing an object* (or stimulus). To express the point very generally, experiences of afterimages typically come without that sense of mind-independence that characterizes normal episodes of perception. Stating precisely what this sense of mind-independence amounts to is not easy, but the following seems right: in perception, we have a sense that we could change our perspective on the objects we perceive, and this sense at least partly derives from the fact that the way in which the appearance of the object changes as we move about is governed by the rules of perceptual constancies – to put it in the same terms as chapter two, objects look to remain the same while looking differently as we move about. By contrast, afterimages *typically* do not to exhibit perceptual constancy, and indeed the phenomenology of experiences of afterimages does not change when we move, or at any rate not in the same way as it does in cases of normal perception.

There are two options for Brewer here: one would be to retract the claim that afterimages are to be treated as hallucination and instead classify them as illusions of some sort. However,

this is not very promising, for two reasons: first, it really seems that there is nothing which one actually sees when one experiences a blue dot after having stared at a red stimulus, so this really seems a case of hallucination; second, if one were to classify it as an illusion for one reason or another, Brewer's account would not have the resources to explain it, as a red dot and a blue dot are not similar to one another. The second option would be to retract or at least revise the epistemic theory of hallucination, and indeed I am inclined to think that afterimages cast a shadow on negative disjunctivism. In either case, Brewer seems to have a problem with physiological illusions such as after images.

Smith thinks Brewer faces serious problems in the case of physical illusion too, although assessing the force of his objections is not straightforward. Let us consider again the case of a green square that looks yellow due to the peculiar lighting. First of all, Smith points out, there is no “relevant” similarity at all between green and yellow. They are simply two quite different colours” (Smith 2010, 393). Second, all the behavioural, noninferential responses of a subject who is unaware of the anomaly in the light conditions “are those appropriate to yellow” (394), not green. Therefore, it is not so clear how the very colour of the square is supposed to constitute one's experience – so the objection goes.

Someone who was sympathetic with Smith's objection might raise analogous perplexities for a purely physical illusion such as that of a straight stick partly submerged in water, for it is not immediately clear how the straight shape of the very stick one sees constitutes one's experience: how is a straight stick similar to a bent stick? They simply are objects with two different shapes. Increasingly, if the subject were to produce behavioural responses purely on the basis of the experience, in cases of physical illusions these would be appropriate to the property that the object is illusorily perceived to have, rather than to those that it actually has. As Smith says, “when I point to the ends of the stick on the basis of how it appears to me, I will not point to their actual location” (394). And this, in turn, makes it much more difficult to understand how the straightness of the stick could figure in one's experience.

It is not immediately clear how to assess the strength of these objections. On the one hand, earlier in this chapter we saw how Brewer strives to accommodate physical illusions by emphasizing the role of the spatiotemporal point of view and the circumstances of perception. On the other, it might be a telling sign that, in connection with physical illusions, he has to include in the notion of visual similarities things that the subject cannot possibly be aware of, for he specifies these as “identities in such things as the way in which light is reflected and

transmitted from the objects in question, and the way in which stimuli are handled by the visual system, given its evolutionary history and our shared training during development” (Brewer 2011: 103). Thus the similarity between a piece of white chalk in red light and a piece of red chalk consists in “the similarity of the light reflected from both” (Brewer 2011: 106), and a straight stick in water is visually similar to a bent stick in the sense that “in the region of space in the vicinity of the eye...light from corresponding parts of the two sticks travels, or would travel, along the same paths” (106).

Smith objects that things that the subject is completely unaware of cannot make anything intelligible for her:

such similarities as these, wholly unknown to the experiencing subject as they typically are, can play no role in an account of illusion that locates error only in an “intelligible” response to a core experience that perfectly matches the actual situation perceived. When I see the yellow-looking square, there is no sense at all in which I see the light travelling to my eyes. The light does not look any way at all to me...nothing I am wholly unaware of can be “taken” by me in any way whatever. (2010: 393)

A lot depends on what is meant here by “intelligibility”. There is of course an understanding of the term on which things such as stimuli and light reflection do indeed make illusion intelligible, by figuring in scientific explanations as to how the environment impacts on our visual system and how the system processes the information received. This is the same sense in which talking of assumptions as being hard-wired in our visual system makes it intelligible how it is that we visually enjoy a three-dimensional world despite our retina receiving only a two-dimensional image that is compatible with an infinite number of three-dimensional configurations; let us call this form of intelligibility *scientific intelligibility*.

However, one might expect Brewer to rely on a different notion of intelligibility, more in tune with the naïve realist project of elucidating the phenomenal character of perception in terms of a relation of acquaintance with mind-independent objects in particular perceptual circumstances; we might call this form of intelligibility *phenomenological intelligibility*. If something like *phenomenological intelligibility* is the *desideratum*, then it is not obvious that Brewer can really make physical illusion intelligible, for one is not conscious of things such as light reflection.

Brewer might reply that Smith's line, together with the objection from phenomenological intelligibility articulated above, are based on a confusion between those factors that determine the way(s) that an object looks on the one hand, and the fact that the object looks those ways on the other: among the factors that determine the looks of an object, we should certainly include things like perceptual processing, retinal images and other things that are not part of the conscious dimension of perception. Those factors, however, determine a conscious fact, namely that an object looks, say, bent. Therefore, conscious acquaintance with the stick just as it is, together with the way in which light is refracted through the water, make it intelligible why the stick looks bent although it is straight – conscious acquaintance and the way the stick looks are genuine phenomenological facts, despite being determined by nonphenomenological factors. It is indeed a general truth about perception that phenomenological facts are grounded in nonphenomenological facts. Suppose one sees an object which is moving away from the place where one is. In the absence of relevant visual cues, it would be unqualifiedly true to say that the object looks smaller and smaller as the distance from the perceiver increases, although in fact it retains its size. Now, it seems natural to say that the phenomenological fact of the object looking smaller is determined by the purely geometrical fact that the image projected by the object on the retina decreases with the distance. Just as in the case of the straight stick that looks bent, here too there is a combination of phenomenological and nonphenomenological factors that contribute to explain why the object looks smaller: given that one is acquainted with the object, and given the relevant facts concerning projective geometry, it is phenomenologically intelligible how the object looks smaller as the distance increases.

I can see only one potential problem with this (hypothetical) response, which has to do with a distinction that Brewer draws between two senses of looks: *thin looks* and *thick looks*. Brewer says:

...o looks F iff o is the direct object of a visual experience from a point of view and in circumstances relative to which o has visually relevant similarities with paradigm exemplars of F. I will say in such cases o thinly looks F. O thickly looks F iff o thinly looks F and the subject recognizes it as an F, or registers its visually relevant similarities with paradigm exemplars of F in an active application of that very concept. (Brewer 2011: 121-22)

While the hypothetical response described above seems very reasonable with respect to thin looks, it is less plausible for thick looks, since these explicitly require registration of the visually relevant similarities that an object has to paradigm exemplars of F. But these visually relevant similarities include things such as light refraction, retinal images and perceptual processing – that is, things that the subject is not in a position to register. Increasingly, explaining thick looks is particularly important, because, as Brewer acknowledges, the “looks locution is...standardly interpreted thickly” (125). So it seems that there still is something problematic with Brewer’s view.

The issue is quite general, and concerns the question whether Brewer can accommodate the presentation of sensible qualities in experience from a purely relationalist standpoint – or, to put it in his own terms, the question of why an object thickly looks F, either veridically or illusorily. Heather Logue has put forward an audacious proposal in defence of Brewer’s account: she thinks that the properties of objects that we perceive are complex ones involving similarities. Commenting on the previously mentioned case of a green square that looks yellow Logue says:

On [Brewer’s] view, the illuded subject doesn’t misperceive the object at all. She fails to perceive its colour, but perceives another property that the object *really has* – viz., the property of being visually relevantly similar to a paradigmatic yellow thing. (A failure to perceive some of an object’s properties doesn’t count as misperceiving it; otherwise, we’d be misperceiving all the time). As it happens, we cannot distinguish colors from these similarity properties, and we have a default tendency to assume that we’re perceiving the former rather than the latter. (Logue 2009: 70)

There are a number of issues that this raises.³⁴ First of all, of there is a genuine question as to why one doesn’t simply perceive the colour of the object, but rather a distinct, colour-related property that the object equally has. As with the appearance-properties view offered by Antony, that is, the problem is whether the view adumbrated by Logue is not just a way of saying that in experience one particular property instead of another is *selected* as *the* property that an object appears to have – which would violate the core purely relationalist commitment to understanding perception simply as presentation with the world.

³⁴ In fact, the view reported by Logue certainly does not coincide with the view upheld by Brewer (2011), although one could see how his (2006, 2008) might have been taken by Logue to suggest that he did. Regardless, the view mentioned by Logue is interesting in itself and thus worth discussing.

Second, if visually relevant similarities are a question of the way light is reflected and stimuli are processed by our visual system, the question still stands how one could perceive those things, since in most cases light doesn't look any way at all to one.

Third, given that an object looking F is a question of similarities both in the illusory and in the normal case, it would seem that we perceive colour-related properties involving similarities, not just colours, in cases of normal perception as well. Thus, my perceiving a yellow banana would be a matter of perceiving not the banana and its yellow colour, but rather the banana and its property of being relevantly similar to a paradigmatic yellow thing; but then it becomes difficult to understand how one could block the unappealing conclusion that all what we ever perceive are colour-related properties involving similarities, not colours.

Logue tries to prevent this last objection with an argument along the lines of Martin's argument against the screening off objection in the context of hallucination, by claiming that the colour-related property in the normal case is only perceived *along with the object's colour*, and that it is the colour which bears all the explanatory weight with respect to the phenomenological and metaphysical properties of the experience:

It's true that yellowness is perceived "alongside" the property of being visually relevantly similar to a paradigm yellow thing. But the explanatory power of perceiving the banana's visually relevant similarity property is "parasitic" on the explanatory power of perceiving the banana's yellowness. If my perceiving the banana's yellowness was an inadequate account of the banana's looking yellow to me, then my perceiving the banana's property of being visually relevantly similar to a paradigm yellow thing would be inadequate as well – but the converse isn't true. Thus, even though both properties are instantiated by the banana, it is plausible that the fact that I perceive the banana's yellowness is doing all the work in accounting for the fact that the banana looks yellow to me. (Logue 2009: 68-69, footnote omitted)

Whatever one may think of Martin's epistemic account of hallucination, I don't think that that strategy easily transposes to the present case, as Logue would have it. The problem is that there isn't any account of how yellowness figures in perception that is distinctive of the veridical case in the first place. For Brewer, perception is fundamentally a relation of acquaintance with particular objects, only in virtue of which those object look certain ways to a subject when the subject registers certain visual similarities. It is obvious that this holds for all cases of perception, regardless of whether they are illusory or not. Therefore, the property

of being yellow cannot play in veridical perception any distinct role that is not already played by the property of being similar to a yellow thing. There is in fact a strong disanalogy with the explanatory order that we find in Martin's epistemic account of hallucination: while in Martin's case we have a positive naïve realist picture of what it is to perceive objects and their features, and an entirely derivative, negative account of what it is to hallucinate something, in Brewer's case there isn't any account of what it is to veridically perceive a property in contradistinction to an account of what it is for an object merely to seem to one to have a certain property; the notion of visually relevant similarities does all the work across the board.

So probably Brewer should not endorse the letter of Logue's proposal, although the idea that perceiving properties is a question of perceiving (registering, noticing) similarities certainly is in tune with the spirit of Brewer's account. In fact, Brewer goes one step farther, for he explicitly endorses a view on which properties *are* a question of similarities. Indeed, he thinks that the truth of a sentence like "*o* is F" "consists in the fact that *o* resembles the F-things; and this is how *o* makes "*o* is F" true (Brewer 2011: 142, fn. 4). Further, he elaborates a nominalism-based analogy between looking F and being F:

Suppose that *o* is F. According to one kind of resemblance nominalism, *o* satisfies the predicate "x is F" in virtue of the fact that *o* sufficiently and appropriately resembles the paradigms whose association with that predicate plays a significant role in determining its meaning. Thus, *o*'s being F is a matter of its resemblance to other things. Still, it is *o* itself that *is* F. Similarly, I claim, if *o* has certain visually relevant similarities with paradigms exemplars of F relative to a given spatiotemporal view and specific circumstances of perception, then *o* itself *looks* F in a perceptual experience of conscious visual acquaintance with that very object from that point of view in those circumstances; (Brewer 2011: 119-120).

Brewer is fond of a version of resemblance nominalism that has been articulated and defended by Rodriguez-Pereyra (2002), on which what is required of an object in order for it to be F is that it be similar to all and only the F-things. Put in terms of truth-makers, this means that what makes it so that a sentence to the effect that *a* is F is true, is the fact that *a* resembles all and only the particulars that are F. It is important to notice that the way in which similarities determine the fact that an object is F is different from the way they determine the fact that it looks F, for an object *being* F is a question of its similarity to *all and only* the F-things (this idea is also known as *egalitarian nominalism*), while an object *looking* F is a question of its

(visually relevant) similarities to *paradigm exemplars* of F-things (reliance on paradigms in defining what it means to be an F-thing is known as *aristocratic nominalism*).

The topic of nominalism is obviously vast, and I am not in a position even to start a proper critical assessment here, so I will simply limit myself to mentioning a couple of things that are relevant to perception. Could a form of resemblance nominalism help Brewer to define the role of sensible qualities in perception? *Prima facie*, it seems useful in explaining how sensible qualities figure in veridical experience: seeing a yellow object, for example, would be a question of seeing an *object* which is yellow, that is, an object that resembles all and only yellow things. Thus the presentation of sensible qualities would consist in the presentation of objects which stand into particular similarity-relations to other objects. A yellow object *looking* yellow under circumstances C, would be a question of there being an object that under C has visually relevant similarities to paradigm exemplars of yellow things.

Things are not so simple, though, for two reasons. First, the relation between the similarity-relations in virtue of which an object is F on the one hand, and the visual similarities in virtue of which an object looks F on the other, is not straightforward, especially in the case of illusion. For an object that is green but looks yellow is *not* similar to yellow things in the sense of similarity that is required for it to be yellow. It has, to be sure, visually relevant similarities to yellow things, where these similarities include things of which the subject is not aware, such as perceptual processing and retinal images. Thus the idea that the actual green colour of an object *o* is presented in a case of illusion in which *o* looks yellow is only as clear as the idea that *o*'s visually relevant similarities to yellow things are presented to a subject; but the subjective presentation of visually relevant similarities was precisely that which was questioned by Smith, so it does not seem that resemblance nominalism is of great help in understanding the role of sensible qualities in perception.

Second, resemblance nominalism might seem problematic anyway, in particular with respect to perception. Rodriguez-Pereyra argues that the properties that a nominalist should seek to explain are lowest determinates (*crimson*, *scarlet*, *vermillion*, say), not determinables (*red*, *blue*, *colour*)³⁵. However, as Bottani (2014) remarks, lowest determinates might “easily be too fine-grained to be perceived, referred and discerned in any way” (7). Further, if determinables are, as Rodriguez-Pereyra says “disjunctions of their determinates” (49), it is

³⁵ The details of the argument do not concern us here, but it is worth saying that it is an argument to the best explanation, where what is to be explained are certain features of the determinate/determinables distinction (see Rodriguez-Pereyra: 49).

not clear that we could perceive determinables either, for, as Bottani points out, “how can one perceives John or Jack, if she perceives neither John nor Jack?” (7). In summary, it is doubtful whether resemblance nominalism is a great game changer for Brewer.

This has been a long journey, and the challenges encountered have contributed in three respects to a qualification of the brand of representationalism which I want to uphold. First, we have seen the necessity of incorporating an element of relationality within a representational view of perception, in the form of singular contents, on pain of “losing the world”. Second, and relatedly, it has become clear that the importance of subjective indiscriminability is relative, as two qualitatively identical experiences might have partly different contents, insofar as they make singular reference to two numerically distinct objects. Third, it has emerged that the representational view I articulate is not opposed to disjunctivism or NR *per se*, but only insofar as (most) proponents of these views hold that one can properly capture the relationality of perception only by dispensing with content.

In the next chapter, I add a different angle to the debate, by considering the question of perceptual content in light of a distinction between particularity and generality in perception. This will bring into relief further inadequacies of purely relationalist views, and in particular of Brewer’s account of looks.

Chapter Four: **Perception, Phenomenal Character and Particularity**

Abstract

This chapter concerns a question that has remained mostly in the background so far: the relation between the phenomenal character and the representational content of perceptual experience. The chapter also gives an assessment of the relationalist attempt at elucidating phenomenal character without mentioning representational content.

More specifically, I do three things: first, I examine an argument for representationalism offered by Siegel and I argue that the argument as it stands is not satisfactory; although I agree with her conclusion and with the most important premise of her argument, I think that much more needs to be done to show that representational content follows naturally from a correct analysis of perceptual consciousness, and the present chapter is indeed an attempt in that direction. Second, I argue that representationalism implies that there is an element of generality in perception, and I show how pure relationalists would emphasize the particularity of perception, instead of its generality. Third, I argue that the grounds for thinking that there is generality in perception is that that better explains its phenomenal and epistemic determinacy. In order to do this, I reconsider Brewer's account of looks in light of the particularity/generality distinction, and I argue that Brewer faces a dilemma: his theory either fails to satisfy important phenomenological and epistemological requirements, or it introduces generality, thus collapsing on a representational view.

Introduction

According to *representationalism*, or *intentionalism*, or the *content view*, perceptual experience is representational; that is, it represents the world as being a certain way (Byrne 2001, Crane 2006, Dretske 1995, Searle 1983, Siegel 2010). According to a different, increasingly popular view, perception is relational, *instead* of representational: it is fundamentally a question of a subject being in a perceptual, non-representational relation to mind-independent objects; call this alternative view *pure relationalism* (Brewer 2006, 2008, 2011, Fish 2009, Travis 2004, 2007, 2013). As we have seen in the previous chapters, the debate invests several dimensions of the philosophy of perception, such as the metaphysics of

perception, its intentionality, its epistemology and its phenomenal character. In this chapter, I am primarily concerned with the topic of phenomenal character, although toward the end I show that what theory of phenomenal character one chooses has far-reaching consequences with respect to epistemological questions.

According to some of the classical representationalist approaches, there is a connection between representational content and phenomenal character – that is, the way perception represents things as being and the way things appear to the subject in perception do not vary independently of one another. Typically, this is expressed as the claim that phenomenal character supervenes on content, meaning that for any two experiences they cannot differ in phenomenal character unless they differ in content between. Thus Byrne states:

...The propositional content of perceptual experiences in a particular modality (for example, vision) *determines* their phenomenal character. In other words: there can be no difference in phenomenal character without a difference in content. (Byrne 2001: 204)

Supervenience claims such as this are often formulated in the context of naturalistic and reductive approaches to consciousness and intentionality in general, and perceptual consciousness and intentionality in particular. At least since the eighties, a number of research programs have dealt with the task of explaining intentionality in terms of representation, and, in turn, representation in terms of a natural relation between internal states of subjects and external states of the world. A supervenience claim is obviously attractive to supporters of those programs, in that it promises an explanation of the seemingly intractable notion of consciousness (or phenomenal character) in terms of intentionality, and then an explanation of intentionality in terms of the naturalistically tractable notion of representation.

Despite the great interest that has surrounded reductionist programs since the groundbreaking works of Dretske (1981) and Millikan (1984), a different way of expressing and conceiving of the relationship between representational content and phenomenal character has recently become quite influential. In particular, there are authors who claim that “phenomenal consciousness has an essential role to play in the theory of intentionality” (Kriegel 2013: 1), in that phenomenal consciousness generates the accuracy conditions that perceptual content is generally equated to. Thus there is a *phenomenal intentionality*: having an experience with, say, “a squarish phenomenal character” when

nothing around you is square, makes the experience “assessable as inaccurate”, while “if the right object or surface *is* square, your experience may be assessable as accurate” (Kriegel 2013, 7; Siewert 1998).¹ According to the phenomenal intentionality approach, the connection between phenomenal character and content is different from the way it has been construed by classical reductive approaches, because, according to phenomenal intentionality theorists, there is at least a kind of content that is “grounded in phenomenal intentionality” (Kriegel 2013: 5), or, perhaps equivalently, is “constitutively determined by phenomenology alone” (Horgan and Tienson 2002: 420)², whereas (what may be called) classic representationalists hold that representational content determines phenomenal character.³

Throughout this work, I have been implicitly leaning toward the phenomenal intentionality approach, at least insofar as my strategy has been to defend representationalism as the best way to make sense of distinctively experiential phenomena such as perceptual constancy (chapter two) and illusion (chapter three), rather than presenting my case for representationalism as a case for reductionist programs. In this chapter, I pursue the same strategy in connection with the phenomenal and epistemic determinacy of perceptual experience.

It is worth pointing out that there are very significant commonalities between the purely relationalist approach that I’ve criticized and the phenomenal intentionality approach to perceptual experience that I feel sympathetic with. In effect, there is at least an important respect in which they are closer to one another than either of them is to a more traditional representational approach to perception, in that they both emphasize *The Significance of Consciousness*, as per Charles Siewert’s (1998) book title. For the standard representationalist, the phenomenal character of perception is after all ornamental, in that it arises from, or supervenes on, the “real thing”, namely natural relations that have come to

¹ Kriegel (2013) traces back the origins of phenomenal intentionality to Brentano and Husserl. Among contemporary philosophers, it is natural to mention Siewert (1998), Horgan and Tienson (2002), and Loar (2003) as those who have contributed the most to reviving the idea. See also Soldati’s claim that “the intentional content of experience is manifest in consciousness” (2013: 462); although Soldati clearly endorses the classical phenomenological idea that phenomenal consciousness plays an essential role in explaining the intentionality of experience, he has qualms about the idea that perceptual experience is *representational* in the sense of having accuracy conditions, so I don’t classify him as a representationalist.

² The two claims are equivalent if we interpret “constitutive determination” as meaning the same as “grounding”.

³ See Byrne’s quote above, but also Dretske 2003b: “According to a representational view of experience, the phenomenal character of our experience is determined not by the objects we experience—these can change or be non-existent while the experience remains, subjectively, exactly the same kind of experience—but by the way the experience represents things to be, the properties it represents objects (if there are any) to have. (69)

acquire definite representational contents over the evolution of organisms endowed with sensory systems. By contrast, both pure relationalists and phenomenal intentionalists conceive of the phenomenal character of experience as having a fundamental explanatory role in elucidating how perception is directed on its objects and discloses the world to our view precisely in virtue of that character. Where lies the disagreement is of course in the particular ways in which phenomenal character is deemed to do so, and the exploration of such ways is one of the main themes of this chapter.

If one compares the more traditional representationalist approach on the one hand, and a purely relationalist account on the other, one might find that the two programs are to a certain extent compatible – to the extent that they are concerned with different objects of research: *perceptual states* for the traditional representational approach, and *perceptual experience* for pure relationalism. In an important representationalist work like that of Burge (2010), notions like phenomenal character or consciousness do not appear frequently, and when they occasionally do, they are carefully kept separated from representation. Indeed, Burge emphasizes that “it is doubtful that any sort of consciousness is necessary for perceptual representation of entities in the environment” (Burge 2010: 188), and that “phenomenal consciousness is not in itself...representational” (402).

Interestingly, pure relationalists might agree with both of Burge’s statements. There should not be much doubt with respect to the second statement, since one of the main claims of pure relationalism is precisely that if perception discloses the world to our view in virtue of its phenomenal character, then it cannot be representational. Things might seem more complicated in the case of the first statement, because pure relationalists think that there is no such thing as perceptual representational content. However, pure relationalists need not deny that there is theoretical validity to those programs of research that explain from a scientific, third-person point of view certain relations between organisms and the environment in terms of representation. All what they need to deny is that the subjective, first-person significance of perceptual experience should be captured in representational terms. Retrospectively, it seems fruitful to read in this light a passage from Travis that I’ve already commented on in chapter two:

Nothing in this makes it illegitimate for cognitive psychologists, in their accounts of the sub-personal processes involved in, say, vision, to speak of representations on which

computations can be performed. My own view is that for such things to serve their explanatory ends, there is no need to see such representations as committed to anything being so (so as *mistaken* or not accordingly). They may simply represent, say, a colour-edge being at a certain position (without representing it as so that it actually is there) (Travis 2004: 59, fn. 3).

Although the passage is far from being crystal clear, I think it can plausibly be interpreted along these lines: one can legitimately talk of one's sensory states as representing their distal causes, without implying that there is, at a personal, experiential level, a commitment contained in the experience to things being the way that they are sub-personally represented to be. The interesting disagreement, then, is not one between those who think that sensory states can be analyzed in representational terms and those who think they can't; maybe this disagreement does not even exist. Rather, the disagreement should be seen as one between those who think, like Kriegel (2013), McDowell (2013) and Siegel (2010), that phenomenal character plays the essential role that it does in disclosing the world to our view in virtue of determining a definite content, and those who think, like Brewer (2011) and Travis (2004, 2013), that phenomenal character can play such a role only if it does not determine by itself a specific content.

Despite there being important differences between the traditional approach to perceptual representation on the one hand, and the phenomenal intentionality approach on the other, they both count as versions of representationalism, in that they both hold that phenomenal character and representational content do not vary independently of one another. By contrast, pure relationalists reject the notion of content, and claim that the appeal to the *objects* of perception is explanatorily fundamental.

In the next section, I discuss an argument given by Susanna Siegel in favor of a representational approach, and I identify a flaw in her defence of the crucial premise. In section two, I argue that what really underlies Siegel's argument, and the notion of perceptual content more generally, is an implicit commitment to the idea that there is generality in perception, and I show how pure relationalists would reject this idea, because they construe perception as wholly particular. More specifically, they think that conceiving of perception as including generality amounts to misconstruing the way in which perception relates us to our environment – the way in which perception is *openness to the world*. In the third section, I examine the prospects for a representationalist response, and I argue that the grounds for

thinking that there is generality in perception is that that better explains its phenomenal and epistemic determinacy. In order to do this, I focus again on Brewer's purely relationalist theory and I argue that it either fails to satisfy important phenomenological and epistemological requirements, or it introduces generality, thus collapsing on a representational view. Incidentally, this will rehabilitate Siegel's argument.

1. A Representationalist Argument

Siegel (2010) argues that there is a connection between phenomenal character, understood in terms of the way things appear to a subject, and representational content. More precisely, she argues that content follows from the simple fact that things appear to be some way to a subject in perceptual experience; insofar as this is her strategy, it seems natural to classify Siegel (2010) as a proponent of the phenomenal intentionality approach (that is, as someone who thinks that there is at least a kind of content that is determined by the phenomenology) although she doesn't use this locution herself. I take the liberty of reconstructing her argument — which she calls *The Argument from Appearing*⁴ — as follows:

P1: "All visual perceptual experiences present clusters of properties as being instantiated" (Siegel 2010, 45);

P2: If an experience E presents a cluster of properties F as being instantiated, then:
things are the way E presents them only if property-cluster F is instantiated;

P3: If things are the way E presents them only if property-cluster F is instantiated, then:

E has a set of accuracy conditions C, such that E is accurate only if property-cluster F is instantiated;

P4: Perceptual contents are sets of accuracy conditions;

Conclusion: All visual perceptual experiences have contents.

I think everyone would agree that the first premise of the argument is crucial, for it makes a substantive claim about the nature of phenomenal character, while all the other premises simply apply to perception in light of independently plausible general principles. To assure

⁴ I take my reconstruction to be faithful to the spirit of the argument, for I simply omit a couple of technical details that would distract our attention from the key issue. A similar argument is offered by Schellenberg (2011). I focus on Siegel's because the way she states the first premise is particularly apt to draw the contrast with purely relationalist views.

ourselves that this is the case, let us briefly discuss premises two to four, and then focus our attention on P1.

Premise two is, I take it, analytically true, for it is simply an application to perception of the following general principle: if a state X presents Y as being the case, then things are the way X presents them only if Y is the case.

Premise three explicitly derives accuracy conditions from the general principle expressed in premise two, and it can also be seen as an instance of a general principle: “if things are the way a state X presents them as being only if conditions C obtain, then X has accuracy conditions that are satisfied in a world only if C obtains” (Siegel 2010: 51).

The conclusion is finally drawn, by means of an implicit equation between content and (a kind of) accuracy conditions⁵, that perceptual experience has content. This equation is a standard move not only in the philosophy of perception, but in the theory of intentionality in general, and it is not likely to raise objections.

Let us then discuss premise one in detail. The main reason why Siegel thinks P1 is correct has to do with structural features of visual phenomenology. She makes her point by means of an analogy between the metaphysics of objects and properties on the one hand, and visual phenomenology on the other:

Consider the claim (made sometimes in discussions of metaphysics) that there is no such thing as a “bare particular”: that is, an object shorn of all its properties. Premise (i) is motivated by the idea that it is not possible for us to represent objects as so shorn in our visual experience...For you to see a cube at all, it must be part of your visual phenomenology that the cube has certain properties: as it might be, having a certain number of facing edges and surfaces, having a certain colour, location, and so on. (Siegel 2010: 46)

Siegel’s reasoning is straightforward: just as there are no bare particulars, there are no instances of perceptual consciousness in which objects are presented as bare particulars either.

⁵ Siegel adds the requirement that accuracy conditions must be conveyed to the experiencing subject. Arguably, the motivation for the requirement lies in the fact that there might be states, mechanisms or processes that have accuracy conditions without having content. A process that regulates the production of a certain molecule in our body, for example, might be supposed to occur only under certain circumstances, and it might be classified as inaccurate, or inappropriate, whenever it occurs under different circumstances. That, however, doesn’t mean that that process has content in the sense in which perceptual experience has content. Content, in the present sense, has to make a difference with respect to the way things are with a subject, for example by guiding her action, prompting her to form appropriate beliefs or by being otherwise manifest in her consciousness. But that can be the case only if the accuracy conditions of the state are conveyed to the subject.

Therefore, some of the properties that objects have must figure in perceptual phenomenology. But if perception presents clusters of properties as being instantiated, then perception has accuracy conditions, and hence content.

Although sympathetic with the conclusion of Siegel's argument, and indeed with its first premise, I think that she oversimplifies in her presentation of the dialectic here, and that there is a sleight of hand in her defence of P1. Much more needs to be done to show that representational content follows naturally from a correct analysis of perceptual consciousness, and the present chapter is indeed an attempt in that direction.

The dialectical landscape considered by Siegel consists of two options, both of which are unpalatable to pure relationalists: either one buys into the conception of phenomenal character which she expresses in P1, or one has to say that we perceive particulars as shorn of their properties. Since pure relationalists would not like the first option, we should consider whether the second is viable; then we will be in a position to see why there is a sleight of hand in Siegel's argument.

Someone who was a supporter of the second option would hold that perception is fundamentally a relation to *objects*, and hence that properties do not figure in perceptual phenomenology. Siegel labels this *radical naïve realism* (RNR):

Radical Naive Realism: All non-hallucinatory experiences consist in a perceptual relation to a worldly item, and properties are not among the things the subject is perceptually related to. (Siegel 2010: 65)

RNR need not deny obvious empirical facts, such as that our visual system registers some of the features of the objects of perception. Rather, its proponents should say that these features are not constituents of the *perceptual* relation which obtains between the subject and the object.⁶ More precisely, RNR seems to be based on the principle that if S is related to an object *o*, then for *any* property of *o* F, it does not follow that S is related to *o*'s F-ness. Thus, if S is perceptually related to a red ball, it does not follow that S is perceptually related to the red colour of the ball.

Even once it is clear that RNR need not be in tension with empirical data, I agree with Siegel that it is phenomenologically implausible, because it makes it difficult to make sense

⁶ On this point see also Logue (2014: 234).

of the fact that the way things appear vary with what one perceives. Being perceptually related to a red cube, for example, is different from being related to a yellow station wagon or a blue cube. However, this difference is made manifest to a subject only in virtue of the different colour, shape and size that these objects appear to have. Indeed, on a plausible, although not uncontroversial view of what it means to see an object *o*, seeing *o*'s features and changes is constitutive of an episode of seeing *o*. Thus Mark Textor writes:

the perception of material particulars is constituted by the perception of their features, states or changes. (Textor 2009: 147)

What would it be to perceive a cube as shorn of all its properties? Would it still look like a cube at all? Since it doesn't even seem possible to make sense of the idea of perceiving bare particulars, then surely it is more promising to say that properties, as well as objects, figure in the phenomenology of visual experience. Before I proceed, I should emphasize that at this stage I am not in a position to offer a conclusive refutation of RNR yet, and that there are things that a proponent of RNR could say in order to explain how the idea of being perceptually related to objects can illuminate how such objects look one way or another. I give full consideration to a possible defence of this radical view in section three, and there I argue that the defence is not sufficient to rescue RNR. For the moment, I simply ask the reader to bear with me.

RNR is not *prima facie* attractive. Does it mean that the representational view is unavoidable? Let us recall Siegel's reasoning: properties have to figure in perceptual phenomenology. But if properties are presented as being instantiated, then there are accuracy conditions and hence contents. This is all a little too quick, though, for there is a sleight of hand from what might be called *middle ground view* to a full-fledged *content view* of the kind that Siegel upholds:

Middle ground view (MGV): properties figure in perceptual phenomenology.

Content view (CV), or *representationalism*: properties figure in perceptual phenomenology *as being instantiated*.

This is not a terminological quibble. The idea expressed by MGv is just that we do not perceive objects as shorn of their properties. That much, however, doesn't quite get Siegel where she wants, since it does not imply that perception attributes properties to particulars, that is, it does not imply that in enjoying perceptual experience we are "told" which properties are instantiated by which objects. MGv is compatible with perception not having content, for accuracy conditions can be generated only if there are properties such that the experience tells the subject, of those properties, that they are instantiated.

How does one get from MGv to representationalism? In order for the argument from appearing to get off the ground, a stronger claim than that properties figure in phenomenology must be made: the claim that visual phenomenology "takes a stand" (Siegel 2010, 47) on which properties are instantiated (and by which objects). That this is what Siegel has in mind is apparent when she writes, in the passage quoted above, that it must be part of the visual phenomenology of an experience of a cube *that* the cube has certain properties (otherwise she might have simply said that certain properties of the cube must be part of phenomenology). MGv only requires that properties be among the *relata* of the perceptual relation, whereas representationalism holds that there are some properties which are presented as being instantiated, and that those properties characterize the way things appear to a subject.

One might still talk of properties figuring in phenomenology as being instantiated while being a supporter of MGv, provided that one used the locution "as being instantiated" only to mean that properties do not appear to float free, isolated from objects. This weak reading of the first premise, however, does not imply content, as it does not imply there being anything in perception which attributes specific properties to objects. According to a stronger reading of the first premise, that properties are presented as being instantiated means that visual phenomenology presents it as being the case *that* things are such-and-such, or *that* certain properties are instantiated – on this stronger reading, "as" indicates not only the fact that properties do not appear to float free, but also that experience takes a stand on which properties are instantiated.

The problem with Siegel's line is that she doesn't give sufficient consideration to the possibility that one might accept MGv (properties figure in perceptual phenomenology) while rejecting representationalism, for she seems to think that the only way in which properties can figure in perceptual experience is by being presented as being instantiated, on the strong reading of the relevant locution. This leads to her presenting her opponent with a false choice:

either buy into her conception of phenomenal character and the inextricably linked notion of content, or end up with the barely understandable claim that we perceive objects as bare particulars. In fact, Siegel would need to provide an argument for getting from the weak to the strong interpretation of P1, but since she conflates them, she is not in a position to do so.

Why would one embrace MGV but reject representationalism, though? If one assumes Siegel's perspective, this must seem inexplicable, for, according to her, the presentation of properties in experience suffices to generate accuracy conditions. But even once the sleight of hand in Siegel's argument is exposed, one might still wonder what is the motivation for MGV. To answer this question, we must conduct a deeper investigation into the most basic grounds for Siegel's argument. Then, we might find something which a pure relationalist would deny, thus blocking the argument from appearing before it even gets started. In the next section, I argue that such basic grounds involve a commitment to the idea that perception includes an element of generality; this, I think, is what pure relationalists would deny.

2. Particularity and Generality

There are several things that might be referred to as *the particularity of perception*, but I have a specific sense of the locution in mind here – one on which one can contrast particular objects, property-instances and events as the objects of perception on the one hand, with what I will call, following Travis (2013) *generalities* as the objects of thought on the other. To anticipate a little: on my interpretation of the debate, the most basic objection that a pure relationalist would raise against representationalism is that it introduces an element of generality into the domain of perception, thereby unduly replacing the objects of perception with objects of thought.

Suppose that John is walking, and that Pia is well placed to see him walking. It would surely be correct to say that Pia sees *John walking*. Would it also be correct to say that Pia sees *that John is walking*, in the sense of perceptually experiencing⁷ *that John is walking*? According to the interpretation that I offer of the debate, a pure relationalist should say that, strictly speaking, Pia can only see *John walking*. The reason why would be, in a nutshell: *John walking* is a particular event, and thus something perceivable. *That John is walking* is a

⁷ That is, not simply in the sense of the locution *seeing that* which is legitimately associated with an epistemic, post-perceptual state.

generality, and thus not something perceivable. There are two important differences between *that John is walking* and *John walking*, the analysis of which will be useful in articulating the distinction between particulars and generalities that I have mentioned. I label the first difference Policy, and the second Invisibility.

Policy: There are a range of things that can count as *that John is walking*. That is, a range of particular configurations of reality that may (or may not) meet a *general* condition on what it is for things to be such that John is walking. It doesn't matter, for example, whether he is wearing a red shirt or a blue one, whereas it might matter whether there are moments at which neither his left nor his right foot touch the ground (for then things would fall under the generality *running*, not *walking*). It is in connection with the particularity/generality distinction that a crucial difference between seeing *that John is walking* and seeing *John walking* emerges: seeing *that John is walking* requires, in contrast to simply seeing *John walking*, a policy for determining whether or not what you see satisfies the general condition on what it is for things to be such that John is walking. However, perception does not provide such policy, because perception is just *openness to the world*.

On the particular conception of openness typically endorsed by pure relationalists, this means that perception simply brings the surroundings into view, putting us in contact with the here-and-now. It does not tell us what generality things fall under. Charles Travis makes this point by saying that it is “no part of what perception is – of how it opens our surroundings to our view – that in perceiving one is to appreciate one set of facts as to what things look like, and ignore others” (2004, 73), and Bill Brewer says that supposing perception to provide a policy for determining whether or not something satisfies a general requirement would be to trade “direct openness to the elements of physical reality themselves, for some intellectual act of classification or categorization” (2006, 174).

Suppose it looks like John is stumbling at every step. For all the perceivable features of what Pia sees, she might categorize the event as one of John dancing, or one of John being drunk, or sick, or one of John practicing a silly walk, as in a famous Monty Python gag. Maybe what she sees is, as a matter of fact, *John practicing a silly walk*, although she gets the impression he's drunk. The fact that she gets a certain impression from what she sees, however, is not sufficient to establish a representational view, for getting a certain impression of John when you see him might be a question of “what you know of him, what you have in mind at the moment, and so on” (Travis 2013: 842). Perception affords Pia awareness of John

while he's practicing his silly walk. However, in order for her to see *that* that is what he's doing, perception should provide her with a policy for determining whether the generality instantiated is one of practicing or one of being drunk, and this is not how perception relates us to the world.

The present point also arises for more basic properties *shape* and *colour*. Suppose you see a circular plate at a skewed angle. Do you see *that* it is circular? A pure relationalist might say no, because that would require a policy for privileging the circular look over the elliptical look that it has from that angle. The same goes for colour: perhaps what you see is a uniformly colored surface that is partly in shadow. But to see *that* the surface is uniform in colour would be to have a policy for determining that the surface satisfies a certain general condition (being uniform in colour under non-uniform illumination), rather than another (actually being of two different colors under uniform illumination)⁸. Generally speaking, for any property F and any kind of event E, to see that a particular object is F, or that a particular event belongs to the kind E, would be to see that the object or the event satisfies "some intelligible demand on membership"(Travis 2007, 231); but *this* is no visible thing.⁹ And this brings me to the next point.

Invisibility: John, and his action of walking, are at a particular location, perhaps just in front of Pia. *That John is walking* is not. Thus Frege states:

That the Sun has risen is not an object emitting rays that reach my eyes; it is not a visible thing like the Sun itself. (Frege 1984: 354)

And thus Travis writes:

⁸ In chapter two, I have given an argument based on constancy for the claim that this purely relationalist way of conceiving of the relation between different looks of an object is in fact wrong. Here I aim at giving an argument for content which does not presuppose my view of constancy though, so I bracket the analysis of looks that I've offered earlier in this work.

⁹ A complementary way of making the same point is in terms of rule-following. In belief, thought and language there are rules which ought to be followed. One should not say, for example, that tomatoes are blue, because the rule for using "blue" forbids it. Indeed, it is constitutive of what it means to think or believe that, say, *a* is F, that one should apply or misapply a rule for the use of a term that is associated with the Fs. What would it be either to obey or violate a rule in perception, though? In perception we are passive, and thus we do not apply or misapply any rule. What the representationalist would have to say in response to this, of course, is that in perception itself there is such an application or misapplication of a general rule, although pure relationalists would retort that there cannot be an application of a rule in perception, because perception is openness to the world.

One sees what is at a particular location...That the sun has set (in Rostock) may be, in some sense, *about* a location. But it has none. The sun, perhaps, is in the sky. That it has set is not. (Travis 2007, 231)

The meat is in the surroundings. To see it, look where it is...To see that the meat is on the rug, you might look where the meat is... You cannot look 'where that the meat is on the rug is'. There is no such place. (2007: 238)

Mark Textor expresses scepticism too:

I can see, for instance, the red surface covering the tomato. But this is not the same as seeing that the tomato is red. What would it mean to see that a particular instantiates a universal? (Textor 2009: 131-132).

Frege, Travis and Textor all express worries about the possibility of propositional seeing, although from slightly different perspectives. Frege's point is related to a metaphysical picture according to which states of affairs are truth-value bearers, and as such abstract entities in the third realm without causal powers. Travis's remark is similar, but with a linguistic twist in it: if something really is an object of perception, you can look for it in your surroundings. If this is true, grammar *suggests* that only particulars are objects of perception, since you can say that you are looking for a particular in your surroundings (looking for the meat), but you cannot say that you are looking for a state of affairs in the surroundings (you cannot look where *that the meat is on the rug is*). Finally, Textor voices general (perhaps phenomenological) puzzlement at the idea of seeing that a particular instantiates a universal.

Despite distinct themes being emphasized in the three passages quoted above, there seems to be a common underlying metaphysical claim to the effect that only particulars are fit to impinge on our sensory system. If we take this metaphysical claim to be the main point of Invisibility, we can use it as the grounds for Policy. Consider a propositional content to the effect that *a* is F. There is a general condition on what it is to be an F, that indefinitely many particular objects may or may not satisfy. To *believe* that *a* is F, for example, would be to believe that *a* has got what it takes to gain access to the club of the Fs. To *see* that *a* is F would be, analogously, to see that *a* satisfies the relevant general condition for being let into the club. But *that* it satisfies the condition is not the kind of thing that can have spatiotemporal location, let alone causally affect our sensory apparatus. A theory of

perception which held that we see that a particular instantiates a universal, then, would place impossible demands on our senses.

Before we move on, it's worth pointing out that the worries I've expounded on propositional seeing are less strictly tied to specific metaphysical outlooks about the constitution of reality than one might suppose. It might be thought, for example, that if there are tropes instead of instantiated universals, the supposed problem of seeing that a particular instantiates a universal simply disappears: seeing that *a* is F would be to see that a bundle of tropes are at a certain spatiotemporal location. However, it certainly seems legitimate to group tropes together in virtue of a certain relation of similarity in which they stand to one another, and talk of a class of tropes, defined as all the tropes that stand in the similarity-relation in question to one another. So Seeing that *a* is F would be a question of seeing that a bundle of tropes belonging to the same class are co-located. But *that* they belong to the same class is not the kind of thing we perceive, thus this brings us immediately back to Policy and Invisibility.

How does all that relate to the question of phenomenal character, though? Consider again Siegel's representationalist conception seen above: she believes that the structure of phenomenal character has to be captured by saying that perceptual experiences present clusters of properties *as* being instantiated. An experience of John walking, for example, would present the individual *John* and the property *walking* as instantiated. This, however, is as general as the propositional content *that John is walking*. For in order for Siegel to be right, perception should tell us which cluster of properties is instantiated, thereby providing the policy to make us see what generality things fall under. However, if the points made above on the behalf of the pure relationalist are correct, we cannot perceptually appreciate generalities.

From a purely relationalist perspective, perception can present you with the property *walking*, if that is what you see, but it cannot present you with the property *walking as* being instantiated. For that, a policy is needed to determine whether or not what is in front of you actually is a case of walking, and perception doesn't provide such policy. The phenomenal character of an experience of John walking thus consists of the experience confronting you with John and his action of walking; not with the fact that John is walking.

We have thus seen one reason why pure relationalists may want to accept MGv but reject CV (representationalism): it is the only way, they would say, to respect the insight that we see particular objects (events, property-instances) in perception. In the next section, I examine the prospects for a representationalist response, and I argue that the grounds for thinking that

there is generality in perception is that that better explains its phenomenal and epistemic determinacy.

3. Perception and Generality

One way of rehabilitating the notion of perceptual content in the face of the objections of section two would be to show that there are independent reasons for thinking that there is generality in perception. In what follows I pursue an indirect strategy, in that I do not make an attempt at showing *how* a policy is given in perception, or give a model of causation on which things with a propositional structure can impinge on our eyes. Rather, I articulate a justification for a refusal to be concerned with these problems: the best way to understand the phenomenal and epistemic nature of mature human perception is to conceive of it such that states of affairs, and not just particulars, can be taken in in perceptual experience.

I think that there are epistemological and phenomenological requirements on any theory of perception, and these have to do with the determinacy of our perceptual experiences:

Phenomenological requirement: a perception of an object *o* determines a suitably narrow range of look statements of the form “*o* looks F (to S)” that would be appropriate to affirm.

Epistemological requirement: a perception of an object *o* determines a suitably narrow range of properties that it would be rational for the perceiving subject to attribute to *o* solely on the basis of her experience.

What “suitably narrow” means here is not a question with a perfectly definite answer, but there are obvious limits on which everyone would agree: if a subject sees a red ball in normal conditions, and a given theory predicts that that ball may look like a white elephant to the experiencing subject, then that theory fails to meet the phenomenological requirement. Likewise, if the theory predicted that a subject to whom an object *o* looks like a red ball would be rational in saying that *o* is a white elephant on the basis of her experience, would fail to meet the epistemological requirement – crucially, less extreme cases would also be relevant: upon seeing two lines that look unequal in length, the subject would not be rational in saying, solely on the basis of her experience, that they are the same length; or, upon seeing

a stick the looks bent, one would not be rational in saying that the stick is straight solely on the basis of the experience. The phenomenological requirement enjoys a certain priority, for what properties would be rational for the subject to ascribe to an object depends on the way that the object in question looks to a subject. Therefore, if a theory were to give odd predictions about the phenomenology, it would give odd predictions about the epistemology too. I think that the best way to satisfy the relevant requirements is to conceive of our perceptual experiences as including an element of generality.

Could a pure relationalist meet the requirements? In section one, I made the promise that I would give more attention to the resources of RNR, and it is now time to keep it. Bill Brewer (2011) has offered a sophisticated defence of a brand of relationalism that is very close to the radical naïve realist rejected by Siegel, in that it holds that perception is fundamentally a relation to an *object* (indeed, Brewer labels his view the *object view*). However, his view also incorporates the insights of the middle ground view, in that it includes an attempt at explaining how properties figure in experience – or, in Brewer’s terms, how, when one sees an object *o*, there are “truths of the form ‘*o* looks *F*’ that qualify in a perfect natural sense to capture how things seem to the subject” (62). Thus Brewer upholds the radical view labelled by Siegel *radical naïve realism*, in that he clearly states that the most fundamental nature of perception consists in a relation to objects, and that properties are not “on a par with the object themselves” (Brewer 2011, 81). At the same time, he has a strategy for avoiding being impaled on Siegel’s objection that the radical view makes it unintelligible why the way things appear vary with what we see. In the present chapter’s terms, it could be said that Brewer’s view is a combination of RNR and MGv:

RNR + MGv: perception is fundamentally a relation to an object, and it is only in virtue of that relation that properties figure in perception.

The question thus becomes whether Brewer can consistently hold a view that satisfies the relevant phenomenological and epistemological *desiderata* on perception on the one hand, but does not collapse on the content view on the other. In the next subsection, I elaborate on these requirements and discuss again Brewer’s account of looks, but this time more explicitly in light of the phenomenological and epistemological requirements, as well as the particularity/generality distinction. In 3.2, I argue that Brewer’s view either fails to satisfy the requirements

or collapses on the content view. Thus the middle ground view, either alone or in combination with RNR, is not a viable option, and representationalism must be endorsed instead.

3.1 *Perception, Determinacy and Similarities*

There are broad epistemological concerns that could be raised both against RNR and MG. Concerning RNR, it is not immediately clear how simply being perceptually related to an object could put one in a position to know specific truths about that object. Hannah Ginsborg, among others,¹⁰ expresses skepticism about the idea that simply seeing an object might rationalize belief:

Perception of an object cannot rationalize belief...unless it presents the object as being a certain way, that is, as having a certain general property or feature. (Ginsborg 2011: 135)

Further, Ginsborg says, even if we broaden the scope of the perceptual relation so as to encompass properties, thus presumably embracing a view along the lines of MG, that does not amount to the relation of inherence between the property and the object being presented to us:

...being presented with an item and with a feature which it has does not yet add up to one's being presented with the kind of connection between the item and the feature which would seem to be needed. (146)

An intuitive way of shaping Ginsborg's worries into an argument against RNR/MG might be as follows: if you see *John walking fast*, it seems that the logical way in which that episode of seeing can rationalize the belief that, say, John is late, or even the noninferential belief that John is walking fast, is by some predicative or proto-predicative relation between John and *walking fast* figuring in your experience. This, in turn, is naturally described by saying that you have to see *that John is walking fast*, which amounts to seeing that there is a certain connection between John and a general predicate that indefinitely many distinct individuals might satisfy. So here is a first, approximate way of articulating the challenge for pure relationalism: if you see John walking, but your senses are silent as to what he is doing (as to

¹⁰ See Logue (2014).

what generality is exemplified), then your being acquainted with that particular event will be silent as well, and therefore epistemically vacuous. Your simply seeing John walking fast does not rationalize any belief (or any belief in particular at any rate), and hence the appreciation of an element of generality (*walking*) in perception seems necessary for it to play an epistemic role.

Pure relationalists would protest that they have the resources to meet the challenge. A pure relationalist such as Mark Kalderon, for example, thinks that perception is a sensory mode of awareness that “is not propositional – it does not take a fact as its object, but a particular” (Kalderon 2011: 220). And yet, Kalderon says, perception has epistemic significance, in virtue of there being an “alethic connection between them [particular objects] and potentially known propositions” (226). It cannot be the case, for example, that I see John walking and the proposition that John is walking is false. Simple perception is epistemically relevant because it makes us aware of the truthmaker for a proposition.

In a similar spirit, Brewer (2011) holds that perception is fundamentally “a relation of acquaintance with certain particulars” (139), and that perception of particulars is epistemically significant in that it acquaints us with “the grounds of empirical truth” (143). So both Brewer and Kalderon think that the contribution that acquaintance makes to perceptual knowledge is to be understood in terms of the connection between the objects that we are acquainted with and certain potentially known propositions that are true in virtue of the way the objects in question are.

However, acquaintance and the alethic connection alone would fall short of providing a satisfactory account of perceptual knowledge, because it would not be clear how the alethic connection can actually be an epistemic connection. Suppose someone sees a red ball and has the true belief that the ball is red popping up in their mind immediately after. Although the object the person sees makes the proposition that the ball is red true, one might doubt that the connection between that episode of seeing and the relevant belief is genuinely epistemic, rather than a merely causal one. In order to have a fully developed epistemic account, the broadly psychological question of how it is that the object in question comes to be appreciated as a reason for acquiring a corresponding belief must be addressed as well.

Brewer is indeed explicit in saying that although acquaintance with a particular is the fundamental *explanans*, perceptual knowledge that *o* is F itself “depends upon far more than mere visual acquaintance with *o*” (Brewer 2011: 144). More precisely, the subject also has (i)

to “*register o*’s visually relevant similarities with paradigm exemplars of F conceptually” (144), and (ii) “actually make the judgment that *o* is F” (145).

The notion of *visually relevant similarities to paradigm exemplars* is absolutely crucial for both the phenomenal and epistemological aspects of Brewer’s theory. According to Brewer’s specific version of pure relationalism, perception is most fundamentally a relation to *objects*, and properties are not “on a par with the object themselves” (81). This immediately raises the question of how Brewer should account for the fact that objects look some way when we perceive them: if objects, but not properties, are among the relata of the perceptual relation, it might seem difficult to account for the fact that object *o* looks, say, F in perception. The notion of visually relevant similarities is in part designed to address this worry. Brewer thinks there are two senses in which an object can look F: it can *thinly* look, F, or *thickly* look F:

...*o* looks F iff *o* is the direct object of a visual experience from a point of view and in circumstances relative to which *o* has visually relevant similarities with paradigm exemplars of F. I will say in such cases *o* *thinly* looks F. *O* *thickly* looks F iff *o* thinly looks F and the subject recognizes it *as an* F, or registers its visually relevant similarities with paradigm exemplars of F in an active application of that very concept. (Brewer 2011: 121-122)

An object looking (either thinly or thickly) F is thus a matter of its being similar to paradigm exemplars of F-things. This raises, in turn, the question of when it is correct to say that *o* has visually relevant similarities to paradigm exemplars of Fs. Let us briefly review how Brewer defines the main elements that make up his theory. He characterizes visually relevant similarities thus:

...visually relevant similarities are identities in such things as the way in which light is reflected and transmitted from the objects in question, and the way in which stimuli are handled by the visual system, given its evolutionary history and our shared training during development. (Brewer 2011: 103)

And here is how paradigm exemplars are defined:

Paradigm exemplars of F-things are “instances of the kinds in question, whose association with the terms for those kinds partially constitutes our understanding of those terms, given our training in the acquisition of the relevant concepts. (104)

Suppose you see a yellow lemon, which looks F (where F is a complex predicate that includes *yellowness*, *sphericality*, *roughness in texture*...). Its looking F is a question of the similarities it has to paradigm exemplars of things that are yellow, spherical, and rough in texture. Although perception is most fundamentally a relation of acquaintance with the *object*, not its properties, it can be made intelligible how it is that that object comes to look F in terms of the similarities, and registration thereof, that the object has with other Fs.

The notion of visual similarities is not simply intended to account for things looking a certain way, though. It also has a central role to play with respect to epistemological concerns. From Brewer’s perspective, being acquainted with an object *o* is a condition whereby we are in a position to actively register *o*’s similarities to other objects and apply a corresponding predicate. If we do so, and the object in question is itself an F, then we attain propositional knowledge to the effect that the object is F. If not, it is still intelligible how we came to mistakenly apply the predicate.¹¹

In the next subsection, I argue that the only way in which Brewer’s account could satisfy the phenomenological and epistemological requirements is by introducing generality in perception, thereby collapsing into a version of representationalism.

3.2 *A Two-Headed Monster and the Generality of Perception*

My main objection to Brewer is, to state it bluntly, that the whole notion of *visually relevant similarities to paradigm exemplars* is a two-headed monster. What is worse, the two heads fight against one another. One head is *visually relevant similarities*, while the other is *paradigm exemplars*. Appreciating the tension within the account, or between the two heads, requires us to investigate the two notions deeper with respect to their nature, their respective function in the theory and how they interact with one another in the explanation of illusion.

Visually relevant similarities are an objective matter: the question whether or not a given object has certain similarities to other objects (given the way the object is, together with

¹¹ In this chapter, I bracket Smith’s worries about the question whether visually relevant similarities are the kind of thing that a subject is in a position to register.

factors such as certain perceptual circumstances, the laws of optics and the constitution of our visual system) is an empirical question with a determinate empirical answer. Further, given the way relevant similarities are defined, any one given object will have similarities with countless other things: a yellow lemon has visually relevant similarities with a yellow lemon, with a golf ball under a yellow light, with an orange under a certain other light, and so on.

The notion of paradigm exemplars, on the other hand, is normative, inasmuch as it is defined in terms of the conditions which are appropriate for the acquisition and intelligible application of an empirical concept. If, for example, you are trying to acquire the concept *equality in length*, and I show you the Müller-Lyer diagram, I am not being very helpful; that does not count as a paradigmatic case of equality in length, in that it is not an instance of equality in length whose association with the expression “equality in length” partly constitutes our understanding of the meaning of that expression. In other words, it is not suitable for acquiring the relevant concept. So the question whether or not some F-thing is a paradigm exemplar of F-things is not purely determined by things such as light reflection and stimuli. It also depends on whether the object can be deemed to satisfy — according to our practices and conventions — a certain general condition on what it means to be F; in short, it depends on normative considerations.

Let us now consider the respective function of the two notions in the theory. *Visually relevant similarities* explains how in perception we are open to the world, on the particular conception of openness that we saw when discussing Policy – one which rules out content. Articulating the notion of visually relevant similarities is a way for Brewer to account for the fact that objects look some way without having to buy into Siegel’s view of phenomenal character, which would indeed contravene to the purely relationalist notion of openness. *Paradigm exemplars*, on the other hand, ensures that the account can respect the phenomenal and epistemic determinacy present in perception. Serious problems would affect Brewer’s theory without the notion of paradigm exemplars. To see these problems one only need to consider illusion and bear in mind the fact that the relation of similarity is symmetrical and reflexive.

Consider the Müller-Lyer diagram (ML): if the lines look unequal in length because the diagram is similar to a diagram where the lines actually are unequal, as Brewer suggests, then the lines in this latter diagram should look equal in length, for the diagram would be similar to the Müller-Lyer diagram, in which the lines are in fact equal in length. Likewise, similarity is

reflexive. However, if the Müller-Lyer diagram is similar to itself, and looks are a question of similarities, then the lines should look equal in length. All this would be unacceptable both from a phenomenological and an epistemological point of view. In particular, it would not respect the previously mentioned requirements: concerning the phenomenological requirement, the theory would predict that the two lines unequal in length which constitute the hypothetical configuration that ML is similar to, would look equal in length, and this seems plainly wrong. From an epistemological point of view, things would not be any better, for the theory would predict that, when the subject is confronted with the two unequal lines in question, it would be rational for her to say that they are the same length.

The notion of paradigm exemplars blocks these unpleasant consequences, because while it is true that plain similarity is symmetrical and reflexive, *similarity to a paradigm* is not: the lines in ML look unequal because the diagram is similar to a *paradigmatic* case of inequality in length, while ML is not a paradigmatic case of equality in length, and therefore the hypothetical pair of lines to which ML is similar is not itself similar to ML. Thus *paradigm exemplars* is absolutely crucial in securing the determinacy that is necessary to exclude odd phenomenological and epistemological predictions.

How does *paradigm exemplars* ensure that there is determinacy, though? In effect, paradigm exemplars constitute a *policy for the selection* of certain similarities, rather than others, as the specific similarities which make it so that something looks the way it looks. However, the driving motivation for pure relationalism which we saw in section two was precisely the need to avoid selection in perception. The challenge for Brewer can be presented in the form of a dilemma. He may either have openness to the world, or phenomenal and epistemic determinacy.

Visually relevant similarities ensures that there is openness, because visually relevant similarities are objectively determined by how an object is, plus the conditions in which it is perceived and other empirical factors. The whole spectrum of similarities that an object has to other objects (given the circumstances of perception) is perceptually manifest to one simply in virtue of seeing the object in question under the relevant circumstances; the object objectively looks (thinly) similar to some other objects. But if Brewer simply avails himself of *visually relevant similarities*, he cannot account for the phenomenological and epistemic determinacy of perception.

If, on the other hand, Brewer avails himself of *paradigm exemplars*, he can have the required determinacy. However, he loses openness to the world. *Paradigm exemplars* introduces an element of selection at a basic level of contact between a subject and the world, for the very notion of an object appearing some way, *even at the level of thin looks* (that is, without conceptual registration), to a subject is partly understood in terms of paradigm exemplars. The role of *paradigm exemplars* is to restrict the otherwise indefinitely extended range of relevant similarities to those similarity-relations that satisfy a certain general requirement: that the objects with which the object of acquaintance has similarities be suitable exemplars for the acquisition and intelligible application of the relevant concept. The very extent to which the similarities that an object has to other objects determine the ways it looks depends on this general element. However, this violates the fundamental relationalist commitment according to which the selective categorization of the fact that *a* is F cannot account for the most basic dimension of perception, which is instead to be understood in terms of conscious acquaintance with the object as it is. Ironically enough, Brewer seems to face the same problem that pure relationalists usually confront representationalists with. Consider again what Travis says on openness:

It is no part of what perception is – of how it opens our surroundings to our view – that in perceiving one is to appreciate one set of facts as to what things look like, and ignore others. (Travis 2004: 73)

And thus Brewer himself writes:

The selective categorization of particular constituents of physical reality enters the picture of a person's relation with the world around her only when questions of their various similarities with, and differences from, other such things somehow become salient in her *thought* about them, rather than constituting an essential part of their subjective presentation in perception. (Brewer 2006: 172)

For Brewer, selective categorization kicks in in thought. However, *paradigm exemplars* aims at explaining *perception*, not thought, and it implies selective categorization in two ways: first, perception must select, from among an indefinitely extended range of similarities, a particular similarity-relation to a paradigm exemplar; second, perception must determine

whether the particular perceived object satisfies a general condition on what it is to be, or to be similar to, the paradigm exemplar in question.

Let us go over the dialectic one last time. We are told that the lines in ML look unequal because the diagram is similar to a diagram constituted of two unequal lines. Now, since similarity is reflexive, it follows that ML is similar to itself, so the two lines should look equal in length. But then, and this is the crucial step, we are told that the reflexive similarity of ML does not count, for it is only similarity to paradigm exemplars that contributes to determining looks. However, this precisely amounts to the claim that a certain fact as to how the diagram looks (that is, its being visually similar to a pair of lines unequal in length), is selected as the fact which makes it so that the lines in ML look unequal. If it were not for the selection of that particular similarity, it would be a mystery why it is that the lines in ML do not look equal to us.

It is apparent that Brewer cannot keep faith to what motivates his view in the first place: *visually relevant similarities* fights against *paradigm exemplars*, because nothing in the way visually relevant similarities are defined leaves room for considerations as to whether an object meets certain general criteria on what a paradigm exemplar is. *Paradigm exemplars* fights against *visually relevant similarities*, because it introduces an element of categorization in the supposedly uncategorized realm of perception.

Following the particularity/generalizity dyad has taken us quite far from the place where we started. Yet, I think seeing the debate in light of that distinction is fruitful in three ways.

First, it makes clear why Siegel's argument can be rejected despite its intuitive plausibility. Siegel's case for P1 is at best incomplete, because she fails to draw a distinction between two ways in which properties might figure in phenomenal character: a) by simply being what the subject is perceptually related to, and b) by being selected as that which an experience represents.

Second, it brings to light a reason why one may want to be a pure relationalist. In order to elucidate perception as a condition whereby one is open to reality, we must conceive of perception (so relationalists say) purely in terms of acquaintance with particulars, for as soon

as we introduce the element of generality that comes with content, we trade direct openness to the elements of physical reality themselves for some intellectual act of classification.

Finally, it indirectly provides the representationalist with the resources to rehabilitate perceptual content. In this connection, I have argued that we need generality in order to make sense of the phenomenal and the epistemic determinacy of perception. With respect to Siegel's account, my line of reasoning removes an obstacle for the first premise of her argument: if generality's being present in perception best explains the epistemic and phenomenal features of our experiences, then there being an implicit commitment to generality in that argument for content should no longer appear problematic. I contend that representationalism is phenomenologically and epistemologically apt.

Conclusion: **Reconciling Naïve Realism and Content?**

Naïve realism (NR) says that perceptual experience is phenomenally and metaphysically constituted by its objects. What does it really mean, and is it compatible with the view that perceptual experience has representational content?

Let's start with the first question. In the first chapter, I described the topic of this work as the question whether perception is openness to the world in the sense of *things being as they are*, or rather in the sense of *that things are thus and so*. Articulating the notion *things being as they are* is a way of capturing the point that in perception we are aware of particular configurations of reality, without it being the case that perception decides what aspects of *things being as they are* matter, and how, to the constitution of a representational content to the effect that things are thus and so. If NR simply meant that perception is awareness of reality in the sense of *things being as they are*, then that would certainly imply an incompatibility with any version of representationalism. Given my arguments for content, I would blame NR.

Maybe there is an alternative way of conceiving NR though, namely as the claim that the phenomenal and epistemological properties of perception, and the way in which perception is directed toward the world, are to be elucidated with reference to its objects and to the essentially relational character that episodes of *seeing* have. This is indeed the way in which most naïve realists, even those who think NR is incompatible with content, formulate the view. Thus Brewer states:

(OV) [Naïve realism] insists that this characterization of perceptual presentation as conscious acquaintance with mind-independent physical objects provides the most fundamental elucidation of which modification of consciousness any specific such experience is: the fundamental nature of perceptual experience is to be given precisely by citing and/or describing those very mind-independent physical objects of acquaintance.
(Brewer 2011: 94)

Brewer talks of elucidating the conscious character of perception with reference to its objects, without mentioning content at all in this particular passage. Could he and other naïve realists stick to this bit, without embracing the overly narrow conception of openness that goes with the rejection of content? I think that this question can be answered in the positive, and that the

view that I have articulated in this work has a good chance of satisfying the explanatory condition placed by Brewer on the elucidation of perceptual experience, while avoiding some of the problems that both he and Travis face due to their rejection of content.

More needs to be said about the sense in which my view could be considered a form of NR. In chapter two (especially in subsection 2.3), I emphasized that there is a two-way explanatory relation between perceptual experience and objects: on the one hand, material objects are revealed as mind-independent only because perceptual experience follows a determinate course, so that the phenomenological constitution of the object as mind-independent is explained by phenomenological features of the experience. On the other hand, it is by reference to mind-independent objects that we understand the actual and counterfactual course of perceptual experience; it is by thinking of a circular glass as something that retains its features over time that we can make sense, as perceiving subjects, of the way it looks differently, without looking to change, as we change our perspective on it.

It will be useful to contrast this approach, on which the object plays an essential explanatory role with respect to the phenomenology of perception, with a fashionable way of elucidating phenomenal character in terms of intrinsic properties of the experiencing subject.

Martine Nida-Rümelin (2011)¹² sets up the problem of perceptual consciousness in the form of a classic inconsistent triad, which I take the liberty of summarizing thus:

(1) When I perceive a crow, the phenomenal character of my experience is fully determined by what I am directly aware of in that perception;

(2) When I hallucinate a crow, my experience has the same phenomenal character as the perception of a crow;

(3) When I hallucinate a crow, there is nothing that I am directly aware of.

On the face of it, the conjunction of (1) and (2) entails not-(3), but it is difficult to see how one could deny (3), since it follows from the very concept of hallucination that one is not aware of an object when hallucinating (sense-data aside, that is). Naïve realists would *typically* deny (2), whereas I take it that most representationalists would rather deny (1). Nida-Rümelin thinks that one can in fact affirm the truth of all of the three propositions above, provided that an ambiguity in the concept of direct awareness is unpacked. She thinks that

¹² Kriegel (2007) adopts a similar strategy for all intentional states.

direct awareness, as is often used, is ambiguous between what is actually *phenomenal presence* on the one hand, and *perceptual awareness* on the other:

...we have to distinguish perceptual awareness and phenomenal presence. Phenomenal character is determined by what is directly present to the subject in the sense of phenomenal presence. The two experiences [the perception and the hallucination of a crow] are not distinct with respect to phenomenal presence; they are only distinct with respect to perceptual awareness.

Perceptual awareness is constituted by causal relations between the object perceived and what is phenomenally present to the perceiving subject. Phenomenal presence, by contrast, is not a relation between the subject and some external object. Phenomenal presence, as I propose to understand it, is not a relation at all. When we talk about what is phenomenally present to a subject we thereby describe the subject's intrinsic, non-relational properties; we do not thereby establish a relation between the subject and something else. (Nida-Rümelin 2011: 353)

According to Nida-Rümelin, perceptual awareness is a causal notion, whereas phenomenal presence is a purely phenomenological concept. What does “direct awareness” mean as used in (1)? If Nida-Rümelin's analysis is right, it must actually mean “phenomenal presence”, for only thus can the veridical perception under consideration be phenomenally the same as a hallucination. By contrast, “direct awareness” as used in (3) must mean perceptual awareness in the relational sense, for in hallucination there is still something that is phenomenally present to one.

Is it all just an equivocation then? Maybe not: there is an immediate problem for Nida-Rümelin, namely that of explaining how phenomenal presence can explain openness to the world in the veridical case. Nida-Rümelin thinks one can answer this by supplementing the account with conditions for linking phenomenal presence to external reality. She identifies three such conditions:

A subject *s* is perceptually aware of an object *o* in a given situation iff for any property *p* of a sufficiently rich class of properties the following holds: (a) *o*'s having *p* is phenomenally present to *s*; (b) the fact that (a) is caused, in the manner typical for perception, by *o*'s having *p*; and (c) relevant changes concerning *o* with respect to this class of properties would lead in that situation—by the normal causal process characteristic of genuine cases of perception—to corresponding changes in what would then be phenomenally present to *s*. (In other words, the following holds for a sufficiently

rich set of properties: would the object cease to have property *p* then the subject would—due to that change and as a causal consequence of that change—cease to be under the impression that there is an object there having property *p*). (Nida-Rümelin 2011: 376, footnote omitted)

According to Nida-Rümelin, there are three conditions which need to obtain in order for a subject to be perceptually aware that an object *o* has *p*: a) that *o* has *p* is phenomenally present to the subject; (b) (a) is caused by *o*'s being *p*; (c) relevant changes in *o* would result in corresponding changes in what is phenomenally present to the subject.

Can this really account for openness to the world? I doubt it. To explain why, it will be useful to pin down what exactly distinguishes this account from a naïve realist outlook. The obvious difference is of course that NR would deny (2), whereas Nida-Rümelin wouldn't. But I think that there is a deeper, more interesting difference. Nida-Rümelin's proposal implies that perceptual awareness is to be *reconstructed* in terms of materials that fall short of the objects of perception themselves, whereas it is the central claim of NR that perceptual awareness is essentially a relation to those objects, in terms of which it has to be elucidated. According to Nida-Rümelin, the subjective dimension of perception is to be understood, at the most basic level, in terms of *intrinsic* experiential properties of the subject, such as *being appeared to thusly*, where these properties do not involve an element of relationality. In order to make sense of the idea that we actually perceive mind-independent objects, she then has to supplement her subjectivist picture with the causal condition that what is phenomenally present be reliably connected to mind-independent reality and changes in it.

There are two related reasons why I think that reconstructing awareness in terms of materials that fall short of objects themselves, as Nida-Rümelin does, is not satisfactory. First, it is not immediately clear why the picture under examination is not a form of indirect realism, on which one is primarily aware of some mind-dependent phenomenally present object, which is in turn causally related to actual mind-independent objects. Second, causality can be plausibly considered an external relation, namely one such that, for any two things that actually are causally related, they could also have been unrelated. A phenomenally present red ball, that is, is causally related to a red ball, but it might equally have been related to a brown chair, had the causal connections been different. But if this is true, this subjectivist picture seems to lose its right to affirm that it is the particular mind-independent red ball that one is presented with, rather than any other object whose presence would causally trigger the same

phenomenally present red ball. John Campbell makes a similar point about Locke's account of ideas and their relation to external reality:

[Ideas] are signs of external phenomena in something like the way in which smoke is a sign of fire. The immediate problem this raises is that although my ideas are signs of their causes, I do not yet know what any of those causes are like. If all I ever get is smoke, how do I know what fire is like? Any causal correlation view will in the end face some version of this question. How can effects provide you, the subject, with any conception of what their causes are like? (Campbell 2002b: 132)

The same challenge applies to the subjectivist picture described above: if all what you are phenomenally presented with is an array of subjective experiential properties, how are you ever going to know what concrete reality is like? It seems that the subjectivist account is not really able to make sense of perceptual experience as openness to the world, as the mind-independent object of perceptual awareness is in fact beyond one's ken, insofar as a merely causal relation between an array of phenomenal properties and its cause does not give a genuine phenomenological insight in the nature of the cause.

My brand of representationalism is different though, for it acknowledges the fundamental role of objects in perception in two ways. First, as mentioned above, I have explained the phenomenal features of experience (its actual and counterfactual course, as Brewer would say) directly with reference to the properties that mind-independent objects of perceptual awareness have – indeed, I have given in chapter two a response to Travis which does not rely on a special phenomenal sense of looks, and is thus compatible with the objectivist insight that looks are something which belong to objects. Second, I have acknowledged in chapter three that perception has singular, object-dependent contents, so that there is always an external element in the individuation of perceptual experiences.

Is then my brand of representationalism actually a version of NR? If NR amounts to the view that perception does not have content, then obviously not. However, I've offered an alternative interpretation of NR, on which the difference between a naïve realist and a non-naïve-realist outlook is that a non-naïve-realist outlook reconstructs perceptual awareness in terms of ingredients, the sum of which actually fail to recreate the epistemic and phenomenological satisfactoriness implied by locutions such as *seeing that p* or *seeing o*. In

this respect, I agree with McDowell when he chastises those broadly Cartesian pictures of experience which are heavily informed by the argument of illusion/hallucination:

So we are trying to reconstruct the epistemic satisfactoriness implicit in the idea of seeing that things are thus and so, using the following materials: first, the fact that it looks to a subject as if things are that way; second, whatever further circumstances are relevant (this depends on the third item); third, the fact that the policy or habit of accepting appearances in such circumstances is endorsed by reason, in its critical function, as reliable. And now the trouble is this: unless reason can come up with policies or habits that will never lead us astray, there is not enough here to add up to what we were trying to reconstruct .
(McDowell 1995: 880)

McDowell's concerns are primarily epistemic, but exactly the same point could be made against an account whose aim was to reconstruct perception as openness to the world in terms of phenomenal presence and causal conditions. McDowell's concerns do not apply to my account though, for the reasons I have mentioned above. To repeat: I think that there is a two-way explanatory relation between perceptual experience and objects: on the one hand, mind-independent objects are phenomenologically constituted as such over the course of a dynamic perceptual experience; on the other, the phenomenal features of experience are explained in terms of mind-independent objects. This suggests a moral: one should start with positing the perceptual relation between the subject and the world, and elucidate the *relata* (the experience on the one hand, and the mind-independent objects on the other) in terms of the characteristics of the relation itself.

So, again, is my brand of representationalism also a kind of NR? Well, I have to conclude by mentioning a problem that I don't feel in a position to solve here: hallucination. In chapter three, I have raised one worry concerning the epistemic account of hallucination, in connection with the treatment of after-images: the objection was that an explanation of after-images in terms of hallucination, although plausible *per se*, implies a commitment to treating them as if they were seeings, whereas in fact they tend to have quite a distinctive phenomenology.

What alternatives are there for a naïve realist account that construes perception as both relational and contentful? One proposal has been put forward by Logue (2013, 2014): according to Logue, NR and representationalism are in fact compatible, provided that one

correctly construes the relations of foundation among seeings, representations and hallucinations. According to Logue's proposal, in the case of veridical perception, we should say that the phenomenal, behavioral and epistemic features of the experience obtain in virtue of its representational content, although the *ultimate* psychological fact in virtue of which the subject represents that things are thus and so is that she perceives (sees) that they are thus and so. In a hallucinatory experience in which it appears that things are thus and so, the relevant phenomenal, epistemic and behavioral features also obtain in virtue of the content of the experience, but here the fact of the subject representing that things are thus and so is the *ultimate* psychological fact in virtue of which the relevant features obtain. In other words: in the veridical case, seeing is bedrock; in the hallucinatory case, representing is bedrock.

A picture like that proposed by Logue would be an ideal resource for those like me who wish to uphold an account on which perceptual experience is both contentful and relational, and I certainly am inclined to endorse it. However, it has to be admitted that Logue's account raises many questions. Although I can't possibly answer the questions here, I would like to mention them as possible topics for further research.

The main question that is left unanswered by Logue is this: what does it mean to say that, in the good case, one represents that things are a certain way in virtue of seeing that they are that way? This actually comprises two questions: a) can we make sense of a representing that obtains in virtue of a seeing, as opposed to a representing that is self-standing, as it were? And b) what does "in virtue of" mean in the present context? The difficulty in answering the first question stems from the fact that it is inherent to the idea of representation that what is represented may or may not exist. Therefore, one might think that whatever goes for the veridical representation goes for the hallucinatory, in the sense that there is no such thing as a particular case of representing depending on a particular episode of seeing. Maybe the relevant difference here would be that between the singular, object-dependent contents of perception on the one hand, and the supposedly existential general contents of hallucination on the other. So the idea would be that singular contents are ultimately grounded in episodes of seeing, whereas existential contents aren't. The only drawback with this way of understanding Logue's proposal is that it discharges most theoretical weight on the answer to the second question, namely what "in virtue of" means.

The second question could in turn plausibly be understood as a question about the relation of *grounding*, namely as the question whether an episode of seeing can *ultimately ground* an

episode of representing. While it might be possible that a specific notion of grounding could do the job, that is a large topic which would need extensive investigation. To sum up: Logue's account would be ideal for combining NR with the thesis that perception is contentful, although as it stands it is only a promissory note.

I have articulated, motivated and defended an account of perceptual experience on which perception being openness to the world is perfectly compatible with its having representational content. Indeed, I've argued that content plays an essential role both in connection with the phenomenological and the epistemological aspects of openness. By contrast, the conception of openness endorsed by Brewer and Travis has been exposed as being too restrictive, and ultimately inadequate to its purpose by their own lights.

References

- Antony, L. (2011) 'The Openness of Illusions', *Philosophical Issues*, 21: 25-44.
- Armstrong, D. (1968) *A Materialist Theory of Mind*. London: Routledge and Kegan Paul.
- Austin, J.L. (1962) *Sense and Sensibilia*. Oxford: Clarendon Press.
- Ayer, A. J. (1956) *The Problem of Knowledge*, London: Macmillan.
- Berto, F. (2013) *Existence as a Real Property. The Ontology of Meinongianism*. (Synthese Library, Vol. 356), Dordrecht: Springer.
- Bottani, A. (2014) 'Two Problems for Resemblance Nominalism', in Reboul, A. (ed.) *Mind, Values and Metaphysics*. Dordrecht: Springer, available online at <http://www.philosophie.ch/kevin/festschrift/Bottani-paper.pdf>.
- Brentano, F. (1995) *Psychology from an Empirical Standpoint*. London: Routledge.
- Brewer, B. (1999) *Perception and Reason*. Oxford: Oxford University Press.
- Brewer, B. (2006) 'Perception and Content', *European Journal of Philosophy*, 14(2): 165-181.
- Brewer, B. (2008) 'How to Account for Illusion', in Haddock, A. and Macpherson, F. (eds.) *Disjunctivism: Perception, Action, Knowledge*. Oxford: Oxford University Press, 168-180.
- Brewer, B. (2011) *Perception and its Objects*. Oxford: Oxford University Press.
- Brewer, B. (2015) 'The Object View of Perception', *Topoi*, DOI 10.1007/s11245-015-9310-y.
- Brewer, B. (unpublished) 'Objects and the Explanation of Perception'.
- Burge, B. (1986) 'Cartesian Error and the Objectivity of Perception' in Pettit, P. and McDowell, J. (eds.) *Subject, Thought and Context*. Oxford: Oxford University Press.
- Burge, T. (1996) 'Our Entitlement to Self-Knowledge', *Proceedings of the Aristotelian Society*, 96: 91-116.
- Burge, T. (2010) *Origins of Objectivity* Oxford: Oxford University Press.
- Byrne, A. (2001) 'Intentionalism Defended' *The Philosophical Review*, 110: 199-240.
- Byrne, A. (2009) 'Experience and Content' *Philosophical Quarterly*, 59: 429-451.
- Byrne, A. and Logue, H. (eds.) (2008) *Disjunctivism*. Cambridge (MA): MIT Press.
- Campbell, J. (2002a) *Reference and Consciousness*. Oxford: Oxford University Press.

- Campbell, J. (2002b) 'Berkeley's Puzzle', in Gendler, T., Hawthorne, J. (eds.), *Conceivability and Possibility*. Oxford: Oxford University Press, 127-144.
- Chisholm, R. (1957) *Perceiving: a Philosophical Study*. Ithaca: Cornell University Press.
- Chudnoff, E. (2012) 'Presentational Phenomenology', in Miguens, S. and Preyer, G. (eds.), *Consciousness and Subjectivity*. Frankfurt: Ontos, 51-72.
- Cohen, J. (forthcoming). 'Perceptual Constancy', in Matthen, M. (ed.) *Oxford Handbook of the Philosophy of Perception*, New York: Oxford University Press.
- Crane, T. (1988) 'The Waterfall Illusion', *Analysis*, 48: 142-47.
- Crane, T. (2001) *Elements of Mind. An Introduction to the Philosophy of Mind*. Oxford: Oxford University Press.
- Crane, T. (2006) 'Is there a Perceptual Relation?', in Gendler, T. S. and Hawthorne, J. (eds.) *Perceptual Experience*, Oxford: Oxford University Press, 126-46.
- Crane, T. (2009) 'Is Perception a Propositional Attitude?', *Philosophical Quarterly*, 59: 452-469.
- Crane, T. (2011) 'The Problem of Perception', *The Stanford Encyclopedia of Philosophy (Spring 2011 Edition)*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/spr2011/entries/perception-problem/>>.
- Crane, T. (2013) *The Objects of Thought*. Oxford: Oxford University Press.
- Davidson, D. (1967) 'The Logical Form of Action Sentences', in Davidson, D. *Essays on Actions and Events*. Oxford: Clarendon Press.
- Davies, M. (1992) 'Perceptual Content and Local Supervenience', *Proceedings of the Aristotelian Society*, 92: 21-45.
- Descartes, R. (1971) *Meditations on First Philosophy*, in *Descartes: Philosophical Writings*, Indianapolis: The Bobbs-Merrill Company.
- Dretske, F. (1969) *Seeing and Knowing*. London: Routledge and Kegan Paul.
- Dretske, F. (1981) *Knowledge and the Flow of Information*, Oxford: Blackwell.
- Dretske, F. (1995) *Naturalizing the Mind*. Cambridge: MIT Press.
- Dretske, F. (2003a) 'Externalism and Self-Knowledge', in Nuccetelli, S. (ed.) *New Essays on Externalism and Self-Knowledge*, Cambridge (MA): MIT Press, 131-142.
- Dretske, F. (2003b) 'Experience as Representation', *Philosophical Issues*, 13(1): 67-82.
- Eilan, N. (1997) 'Objectivity and the Perspective of Consciousness', *European Journal of Philosophy*, 5: 235-50.

- Evans, G. (1982) *The Varieties of Reference*. Oxford: Clarendon Press.
- Fish, W. (2009) *Perception, Hallucination and Illusion*. New York: Oxford University Press.
- Fodor, J. (1987) *Psychosemantics*. Cambridge (MA): MIT Press.
- Foster, J. (2000) *The Nature of Perception*. New York: Oxford University Press.
- Frege, G. (1984) *Collected Papers on Mathematics, Logic and Philosophy*. Oxford: Blackwell.
- Ginsborg, H. (2011) 'Perception, Generality and Reason', in Reisner, A. and Steglisch-Petersen, A. (eds.) *Reasons for Belief*. Cambridge: Cambridge University Press, 131-157.
- Gregory, R. (1997) 'Visual Illusions Classified', *Trends in Cognitive Sciences*, 1(5): 190-194.
- Grice, H. P. (1961). 'The causal theory of perception, part I', *Proceedings of the Aristotelian Society*, 121: 121-152.
- Gupta, A. (2006) 'Experience and Knowledge', in Gendler, T.S and Hawthorne, J. (eds.) *Perceptual Experience*, Oxford: Oxford University Press, 181-204.
- Haddock, A. and Macpherson, F. (2008) 'Introduction: Varieties of Disjunctivism', in Haddock, A. and Macpherson, F. (eds.) *Disjunctivism: Perception, Action, Knowledge*. Oxford: Oxford University Press, 1-24.
- Hawley, K. and Macpherson, F. (eds.) (2011) *The Admissible Contents of Experience*. Wiley-Blackwell.
- Higginbotham, J. (2002) 'On Events in Linguistic Semantics', in Higginbotham, J. Pianesi, F. and Varzi, A. (eds.) *Speaking of Events*, New York: Oxford University Press.
- Hopp, W. (2011) *Perception and Knowledge. A Phenomenological Account*. Cambridge: Cambridge University Press.
- Horgan, T. and Tienson, J. (2002) 'The Intentionality of Phenomenology and the Phenomenology of Intentionality', in Chalmers, D. (ed.), *Philosophy of Mind: Classical and Contemporary Readings*. Oxford : Oxford University Press.
- Huemer, M. (2011) "Sense-Data", *The Stanford Encyclopedia of Philosophy* (Spring 2011 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/spr2011/entries/sense-data/>>.
- Hume, D. (1960) *Treatise on Human Nature* London: Oxford University Press.
- Husserl, E. (2001) *Logical Investigations*. London: Routledge.

- Jackson, F. (1977) *Perception*. Cambridge: Cambridge University Press.
- Kalderon, M. (2011) 'Before the Law', *Philosophical Issues*, 21: 219-244.
- Kelly, S. (2010) 'The Normative Nature of Perceptual Experience', in Nanay, B. (ed.) *Perceiving the World*. New York: Oxford University Press, pp. 146-159.
- Kriegel, U. (2007) 'Intentional Inexistence and Phenomenal Intentionality', *Philosophical Perspectives*, 21: 307-340.
- Kriegel, U. (2013) 'The Phenomenal intentionality Research Program', in Kriegel, U. (ed.) *Phenomenal Intentionality*. New York: Oxford University Press, 1-26.
- Loar, B. (2003) 'Phenomenal Intentionality as the Basis for Mental Content'. In Hahn, M., Ramberg, B. (eds.), *Reflections and Replies: Essays on the Philosophy of Tyler Burge*. Cambridge (MA) : MIT Press.
- Logue, H. (2009) *Perceptual Experience: Relations and Representations*. PhD Thesis, MIT.
- Logue, H. (2012) 'Why Naïve Realism?', *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 Brogaard, B. (ed.) *Does Perception Have Content?* Oxford: Oxford University Press, 220-240.
- Martin, M (1994) 'Perceptual Content', in Guttenplan, S. (ed.) *A Companion to the Philosophy of Mind*. Oxford: Blackwell.
- Martin, M. (1997) 'The Reality of Appearances'. In Sainsbury, M. (ed.) *Thought and Ontology*. Milan: Franco Angeli. Reprinted in Byrne, A. and Logue, H. (eds.) (2008), *Disjunctivism: Contemporary Readings*. Cambridge (MA): MIT Press, 91-116.
- Martin, M. (2002) 'The Transparency Of Experience', *Mind and Language*, 17(4): 376-425.
- Martin, M. (2004) 'The Limits of Self-Awareness', *Philosophical Studies*, 120: 37-89.
- Martin, M. (2006) 'On being alienated', in Gendler, T. and Hawthorne, J. *Perceptual Experience*. Oxford: Oxford University Press.
- Martin, M. (2010) 'What's in a Look', in Bence, N. (ed.) *Perceiving the World*. New York: Oxford University Press, 160-225.

- Matthen, M. (2010) 'How Things Look (and What Things Look That Way)', in Nanay, B. (ed.) *Perceiving the World*. New York: Oxford University Press, 226-51.
- McDowell, J. (1982) 'Criteria, Defeasibility and Knowledge', *Proceedings of the British Academy*, 68: 455-479.
- McDowell, J. (1984) 'De Re Senses', *The Philosophical Quarterly*, 34(136): 283-294
- McDowell, J. (1994) *Mind and World*. London: Harvard University Press.
- McDowell, J. (1995) 'Knowledge and the Internal', *Philosophy and Phenomenological Research*, 55: 877-93.
- McDowell, J. (2008) 'Responses', in Lindgaard, J. (ed.) *Experience, Norm, and Nature*. Oxford: Blackwell.
- McDowell, J. (2009) *Having the World in View*. Cambridge: Harvard University Press
- McDowell, J. (2013) 'Perceptual Experience: both Relational and Contentful', *The European Journal of Philosophy*, 21(1): 144-157.
- McGinn, C. (1982) *The Character of Mind*. Oxford: Oxford University Press.
- McKinsey, M. (1991) 'Anti-Individualism and Privileged Access', *Analysis*, 51(1): 9-16.
- Millikan, R. (1984) *Language, Thought and Other Biological Categories*, Cambridge (Mass.): MIT Press.
- Nida-Rümelin, M. (2011) 'Phenomenal presence and Perceptual Awareness: a Subjectivist Account of Perceptual Openness to the World', *Philosophical Issues*, 21: 352-383.
- Nudds, M. (2009) 'Recent Work in Perception: Naïve Realism and its Opponents', *Analysis*, 69(2): 334-346.
- O'Shaughnessy, B. (2003) *Consciousness and the World*. Oxford: Oxford University Press.
- Palmer, S. (1999) *Vision Science: Photons to Phenomenology*. Cambridge (MA): MIT Press.
- Pautz, A. (unpublished) The Object View and the Content View: Comments on Brewer, Perception and its Objects. APA Author Meets Critics, Chicago.
- Price, H. (1932) *Perception*. London: Methuen.
- Pritchard, D. (2012) *Epistemological Disjunctivism*, Oxford: Oxford University Press.
- Robinson, H. (1994) *Perception*. London: Routledge.
- Rodriguez-Pereyra, G. (2002) *Resemblance Nominalism*. Oxford: Oxford University Press.

- Russell, B. (1912) *Problems of Philosophy*. London: Thornton Butterworth.
- Salmon, N. (1989) 'Illogical Belief', *Philosophical Perspectives*, 3: 243-285.
- Schellenberg, S. (2008) 'The Situation-Dependency of Perception', *The Journal of Philosophy*, 105(02): 55-84.
- Schellenberg, S. (2011) 'Perceptual Content Defended' *Nous*, 45(4): 714-750.
- Searle, J. (1983) *Intentionality*. Cambridge, Cambridge University Press.
- Siegel, S. (2004) 'Indiscriminability and the Phenomenal', *Philosophical Studies*, 120: 91-112.
- Siegel, S. (2005) 'The Contents of Perception', *The Stanford Encyclopedia of Philosophy* (Spring 2006 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/entries/perception-contents/>.
- Siegel, S. (2010) *The Contents of Visual Experience*. Oxford: Oxford University Press.
- Siegel, S. (2013) 'Replies to Campbell, Prinz, and Travis', *Philosophical Studies*, 163 (3): 847-865.
- Siewert, C. (1998) *The Significance of Consciousness*. Princeton: Princeton University Press.
- Siewert, C. (2007) 'In Favor of (Plain) Phenomenology', *Phenomenology and the Cognitive Sciences*, 6: 201-220.
- Shoemaker, S. (1994) 'Phenomenal Character', *Nous*, 28(1): 21-38.
- Smith, A. D. (2002) *The Problem of Perception*. Cambridge (MA): Harvard University Press.
- Smith, A. D. (2010) 'Disjunctivism and Illusion', *Philosophy and Phenomenological Research*, 80(2): 384-410.
- Soldati, G. (2012) 'Direct Realism and Immediate Justification', *Proceedings of the Aristotelian Society*, 122(1): 29-44.
- Soldati, G. (2013) 'Elements of a Phenomenological Theory of Perception', *Rivista di filosofia*, 104(3): 461-484.
- Sollberger, M. (2013) 'In Defence of a Structural Account of Indirect Realism', *European Journal of Philosophy*, DOI: 10.1111/ejop.12015.
- Soteriou, M. (2000) 'The Particularity of Visual Perception' *European Journal of Philosophy*, 8(2): 173-189.

- Soteriou, M. (2009) 'The Disjunctive Theory of Perception', *The Stanford Encyclopedia of Philosophy* (Fall 2009 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2009/entries/perception-disjunctive/>>.
- Snowdon, P. (1992) 'How to Interpret "Direct Perception"', in Crane, T. (ed.) *The Contents of Perception*. Cambridge: Cambridge University Press, 48-78.
- Strawson, P. (1959) *Individuals*. London: Methuen.
- Strawson, P. (1966) *The Bounds of Sense*. London: Methuen.
- Strawson, P. (1979) 'Perception and its Objects', in: MacDonald, G. (ed.) *Perception and Identity: Essays Presented to A. J. Ayer*. Macmillan, London
- Textor, M. (2009) 'Are Particular or States of Affairs Given in Perception?' in Reicher, M. E. (ed.) *States of Affairs*. Frankfurt: Ontos, 127-150.
- Travis, C. (2004) 'The Silence of the Senses', *Mind*, 113: 57-94.
- Travis, C. (2005) 'A Sense of Occasion', *Philosophical Quarterly*, 55(219): 286–314.
- Travis, C. (2007) 'Reason's Reach', *European Journal of Philosophy*, 15(2): 225-248.
- Travis, C. (2008) 'Pragmatics', in Travis, C. *Occasion-Sensitivity. Selected Essays*. Oxford: Oxford University Press, 109-129.
- Travis, C. (2013) 'Susanna Siegel, the Contents of Visual Experience', *Philosophical Studies*, 163 (3): 837-846.
- Williamson, T. (1990) *Identity and Discrimination*. Oxford: Blackwell.
- Williamson, T. (2000) *Knowledge and its Limits*. Oxford: Oxford University Press.
- Zahavi, D. (2003) *Husserl's Phenomenology*. Stanford: Stanford University Press.