

9-2018

Theories of Perception and Recent Empirical Work

Philip Zigman

The Graduate Center, City University of New York

[How does access to this work benefit you? Let us know!](#)

Follow this and additional works at: https://academicworks.cuny.edu/gc_etds

 Part of the [Philosophy of Mind Commons](#)

Recommended Citation

Zigman, Philip, "Theories of Perception and Recent Empirical Work" (2018). *CUNY Academic Works*.
https://academicworks.cuny.edu/gc_etds/2772

This Dissertation is brought to you by CUNY Academic Works. It has been accepted for inclusion in All Dissertations, Theses, and Capstone Projects by an authorized administrator of CUNY Academic Works. For more information, please contact deposit@gc.cuny.edu.

THEORIES OF PERCEPTION AND RECENT EMPIRICAL WORK

by

PHILIP ZIGMAN

A dissertation submitted to the Graduate Faculty in Philosophy in partial fulfillment of the requirements for the degree of Doctor of Philosophy, The City University of New York

2018

© 2018

PHILIP ZIGMAN

All Rights Reserved

Theories of Perception and Recent Empirical Work

by

Philip Zigman

This manuscript has been read and accepted for the Graduate Faculty in Philosophy in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

David Papineau

Date

Chair of Examining Committee

Nickolas Pappas

Date

Executive Officer

Supervisory Committee:

David Rosenthal

David Papineau

Howard Robinson

Jesse Prinz

THE CITY UNIVERSITY OF NEW YORK

Abstract

Theories of Perception and Recent Empirical Work

by

Philip Zigman

Advisor: David Rosenthal

In this dissertation I answer the following question: Does recent empirical work give us reason to think that naïve realism is false or that indirect realism is correct? There is a small amount of literature arguing that recent empirical findings pose problems for naïve realism and suggest that perception involves mental representation. I review this literature and the arguments therein, examine the relevant empirical work, and argue that recent empirical work on perception does give us reason to reject naïve realism and to favour an indirect realist view that countenances mental representations.

Acknowledgments

I want to thank Howard Robinson, David Papineau, and especially David Rosenthal, for their guidance and support. I also want to thank Jesse Prinz for graciously agreeing to come on board with unreasonably short notice.

And I want to thank Robyn for everything, but mostly the hugs.

Contents

Introduction	1
Chapter 1 - Perception and Representation	7
Chapter 2 - Recent Empirical Work and the Nature of Perception: A Review of the Literature	58
Chapter 3 - Two Arguments Against Naïve Realism	102
Chapter 4 - Concluding Remarks: Veridicality, Prediction, and Implications	171
Bibliography	206

Introduction

Contemporary scientifically minded discussions of perception often note that perception is active or that it is construction. It is equally common to come across talk of our brains deciding what we see, or to find proclamations about how we do not really experience what is out there in the mind-independent world. Here are just a handful of examples:

Instead, this observation suggests that the filled-in percept is actively constructed by neural processes in the brain. (Komatsu, 2006, p. 223)

In daily life, the constructive and creative nature of visual perception becomes apparent when visual information is only partially available because of occlusion. (Wokke et al., 2013, p. 63)

...our perception of the world [is] a construction that does not accurately represent the outside... (Eagleman, 2011, p. 28).

Thus what we see does not depend entirely on what is out there but also to a considerable extent on what the brain computes to be most probably out there. (Smythies, 2005, p. 18)

Modern neuroscience has demonstrated that the content of our conscious experience is not only an internal construct but also an extremely selective way of representing information. ...[O]ur brains generate a world-simulation, so perfect that we do not recognize it as an image in our minds. (Metzinger, 2009, p. 6-7)

Our brains build models of the world and continuously modify these models on the basis of the signals that reach our senses. So, what we actually perceive are our brain's models of the world. (Frith, 2007, p. 134)

Such claims would seem to have important consequences for how we understand perception. And if recent empirical findings are inspiring these sorts of claims, one can hardly be blamed for thinking that this empirical work can potentially shed light on how our perceptual experience is related to the mind-independent world.

In this dissertation I look at recent empirical work on perception and consider whether empirical findings indeed give us reason to favour certain theories of perception over others. The theories of perception I focus on concern the nature of perception, so how we should fundamentally understand and explain perception. This involves articulating just how perception relates us to reality, as well as explaining the phenomenology of experience.

When we have a perceptual experience things look a certain way, there are sounds and smells, tastes and sensations of touch—there are different sensory qualities that appear to us, which we can describe, and on the basis of which we can compare and contrast different experiences. We may sometimes talk about how things seem to oneself, or what an experience is like, or how an experience feels. In doing so we describe the phenomenal character of that experience. My perceptual experience as I sit at my computer today has a very similar phenomenal character to my perceptual experience as I sat at my computer yesterday. And what my experience is like would change if I closed my eyes, or put on some music, or if someone poured a bucket of water on my head.

Different theories about the nature of perception may explain the phenomenal character of our experience in different ways. This is not terribly surprising. If perception is a matter of our being directly related to mind-independent objects and events, this will shape how we explain the phenomenal character of perceptual experience. The same is true if perception is a matter of our being related to certain mind-dependent objects, or if perception is a matter of our representing

the world. A theory of perception may also have far-reaching consequences regarding a variety of topics in philosophy, ranging from epistemological issues concerning our beliefs about the world, metaphysical issues about the nature of reality, and various issues in the philosophy of mind. Here, I restrict my focus to the basic question of how we should fundamentally understand perception, as well as the closely related issue of explaining the phenomenal character of our perceptual experience.

At the heart of this dissertation is the following question: Does recent empirical work give us reason to think that naïve realism is false or that indirect realism is correct? There is a small amount of literature in which recent empirical findings are marshalled in opposition or support of one or another theory about the nature of perception. This project follows in that vein of looking at whether these empirical findings can help adjudicate the debate between different philosophical theories. I argue that certain findings—concerning binocular rivalry, multimodal perception, temporal filling-in, visual completion, spatial perception, colour perception, and speech perception—give us reason to reject naïve realism and to favour an indirect realist view that countenances mental representations.

I will begin the first chapter by surveying salient contemporary theories about the nature of perception. I restrict my focus to realist views. Not only are such views dominant today, but given that I am interested in whether scientific work can inform the debate, considering views which deny the existence of a mind-independent physical world seems inappropriate. After examining some of the arguments motivating these different approaches, I will argue that we should think of views that accept the common kind assumption—i.e. that take veridical experiences, illusions, and hallucinations to have the same fundamental nature—as being views on which perception is in an important sense indirect. I then argue that views which take

perception to involve mentally representing the world are indirect in the same way; and, that focusing on the issue of mental representation instead of the common kind assumption is advantageous. I conclude the first chapter by arguing that the opposition between naïve realism and indirect realism is particularly fruitful to consider. I will argue that we can understand the two views as standing on opposite sides of the question of whether perception involves mental representation, so if empirical work can help us answer this question it will resolve the debate between the philosophical theories.

In the second chapter I will review the literature in which it is argued that recent empirical work gives us reason to think that naïve realism is false or that indirect realism is correct. First, I look at Nanay's (2014) argument that the processing involved in multimodal perception requires us to countenance mental representations. Second, I examine Smythies and Ramachandran's (1997) argument that certain findings—regarding a binocular rivalry study with patchwork images and various kinds of filling-in—indicate that the character of our experience is determined by a representation generated by the brain, rather than by external objects. Third, I examine Brown's (2008) argument for indirect realism, which he considers a new argument from illusion based on empirical findings regarding illusory contours and colour perception. Fourth, I consider Pautz's (2016) argument that we should reject naïve realism because the similarity relations among sensible properties are poorly correlated with the similarity relations among the objective properties of objects. I will argue that the most promising argument is the one offered by Smythies and Ramachandran. This argument, which I call “the mismatch argument,” focuses on cases where there is a mismatch between what our perceptual experience is like and the actual objects and events before us in the world.

In the third chapter I will defend the mismatch argument and a second argument, which I call “the argument that mismatches are ubiquitous.” Both arguments target naïve realism and support indirect realism. I will begin the chapter by closely examining empirical work relating to mismatches. I discuss findings concerning binocular rivalry, multimodal perception, temporal filling-in, visual completion, spatial perception, colour perception, and speech perception. This discussion will make plain that rather than being exceptional occurrences elicited in the confines of the lab, mismatches—often slight and inconsequential—are the norm in our day-to-day lives. I will then support the three key claims found in the mismatch argument and the argument that mismatches are ubiquitous. First, I argue that we can best explain our experiences in mismatch cases by invoking mental representations. Second, I argue that perceptual processing always operates in the same way, so if the brain generates a representation in some cases it does so in all cases. Third, I argue that the ubiquity of mismatches is ruinous for naïve realism. I conclude that recent empirical work gives us reason to think that perception involves mentally representing the world, such that the phenomenal character of our perceptual experience is determined by inner representations.

In the concluding chapter I will briefly explore the issue of perceptual veridicality in light of the fact that mismatches are ubiquitous. I will argue we should adopt a looser notion of veridical perception, as lots of normal perception we are right to consider perfectly good will include minor mismatches.

One issue I will not explore is whether unconscious perception is an insurmountable problem for naïve realism, as has been argued (Nanay, 2014; Berger and Nanay, 2016). Although I think there is substantial empirical support for unconscious perception—whether from masked priming studies or phenomena like blindsight—the issue is somewhat controversial within

philosophical circles (see, e.g., Phillips and Block, 2016; Peters et al., 2017; Phillips, forthcoming). For this reason, and because I can answer the question at the heart of my project despite doing so, I will avoid this issue. Note, though, that there is some debate about whether unconscious perception in fact poses a problem for naïve realism (Zieba, 2017; Anaya and Clarke, 2017).

Chapter 1

Perception and Representation

1. Introduction

Much like being alive, perceptual experience is at once pretty humdrum and fairly incredible. As we go about our lives we see and hear and smell and taste and feel things. In a way, that is just what it is to be awake and get on in the world. We have perceptual experiences as we navigate our environment and interact with objects and other people, as we learn things about the world, as we form memories that we can later reflect on or relive. Perceptual experience seems simply the medium of life.

Yet, there are some perceptual experiences that may stand out from the rest. There are certain tastes that are so enjoyable we will seek them out, particular sounds that move or delight us. We climb mountains and hike through forests to take in natural scenes, cross oceans to stare at paintings in a room.

And even if we ignore the perceptual experiences we may pursue and consider the most mundane, perceptual experience may still seem astonishing and puzzling. That we have perceptual experience at all can seem striking. That our perceptual experiences are a certain way, as opposed to being some other way, may seem curious. And how our perceptual experiences arise, as well as what they may tell us about ourselves and our world, can come to seem like mysteries worth investigating.

Take the relatively mundane experience of my sitting at my computer, my forearms resting on the table, the sound of trees blowing in the wind coming in through the window, while I stare attentively at the computer screen. I can easily describe how things seem to me, or what

appears before me. I can describe the shape of my computer, the colours on the screen, the lamp to the right, the feeling of pressure below my elbows, the intermittent rustling. At a certain level it is obvious just what is going on: I am sitting at my computer, which I can see and feel in front of me, I am sitting on the chair I can feel beneath me, and so on. But it also seems fair to wonder about what is going on in terms of myself and my mind, or even in the world, such that I have this perceptual experience and that it has the characteristics it does. Fortunately for the curious, philosophers have been keen to understand perception, often in the context of a more general understanding of the mind and the world. These efforts have provided no shortage of theories.

Here, I will begin by surveying salient contemporary views about the nature of perception. I will then look at different considerations motivating these approaches and by which we can distinguish the views. Regarding the crucial question of whether perception involves mentally representing the world, I will argue that naïve realism stands in opposition to the other views surveyed, which can all be understood as similarly indirect. Finally, I will argue that the opposition between naïve realism and indirect realism is particularly fruitful to consider.

2. Theories of Perception

2.1. Indirect Realism

According to indirect realism, perceptual experience involves the awareness of mind-dependent objects, which are themselves caused by mind-independent objects and events. We perceive mind-independent objects when we are aware of the mind-dependent objects they cause and which represent them. But while we perceive mind-independent objects we are only indirectly aware of them. What we are directly aware of in experience are mind-dependent objects and the phenomenal qualities that belong to them. That is just to say that how things

appear to us in experience is due to the internal objects. Or, in the modern vernacular, the phenomenal character of experience is determined by the internal objects.

On such a view, when I sit at my table looking at my computer I see my computer, which is an external object. But the sensory qualities that appear to me, that make up my experience—the colours and shapes and sounds—are mind-dependent.

Some versions of indirect realism construe experience in terms of sense-data. Such theories were quite a bit more popular in the first half of the twentieth century (e.g. Russell, 1912; Broad, 1923), though more recent accounts have been offered (Jackson, 1977; Robinson, 1994).¹ On sense-data theories, sense-data are the immediate objects of awareness. Thus, when I experience the silver rectangular border around my computer screen, I am aware of a sense-datum with a silver rectangular border. Russell construes perceptual experience in this way:

Let us give the name of “sense-data” to the things that are immediately known in sensation: such things as colours, sounds, smells, hardnesses, roughnesses, and so on. We shall give the name “sensation” to the experience of being immediately aware of these things. Thus, whenever we see a colour, we have a sensation *of* the colour, but the colour itself is a sense-datum, not a sensation. The colour is that *of* which we are immediately aware, and the awareness itself is the sensation. (p. 12)

Some defenders of sense-data theory claim that we perceive sense-data. When it comes to visual experience the claim is that we see sense-data. According to such views, we directly perceive private sense-data and only indirectly perceive external objects. It is not simply claimed that we are indirectly aware of external objects via being aware of sense-data, but that sense-data

¹ Of course, not all sense-data theorists are indirect realists. Some defend phenomenalist or idealist views (see, e.g., Ayer, 1940; Foster, 2000). But, as noted in the Introduction, I will restrict my focus to realist views.

are the direct objects of perception and play a mediating role in our perception of external objects. Austin (1962) took issue with “directly” and “indirectly” perceiving as described by such theories, arguing, among other things, that the uses of such expressions ran afoul of any normal use. And, indeed, it is unclear whether claims of directly perceiving sense-data do any theoretical work, or if the issue is simply linguistic. If two theorists agree that in perceptual experience we are aware of sense-data and that sense-data are caused by mind-independent objects, then disagreement about whether or not we see or perceive sense-data seems like a debate about the meaning of those expressions. It is not clear what is added to a theory by the claim that sense-data are the objects one perceives. If there is no theoretical difference, it seems sensible to drop the claim and avoid debates about the analysis of “perceive,” “see,” and related expressions. However, whether or not the issue is substantive, I think it is at least misleading to say that we perceive sense-data, so will avoid doing so.

Not all versions of indirect realism construe experience in terms of sense-data. The mind-dependent things we are aware of in experience are often called “impressions,” “sensa,” “qualia,” or “ideas.”² Some recent accounts opt for such alternatives (Brown, 2008; Sollberger, 2015).

Sollberger defends an indirect realist theory on which mind-independent objects cause mind-dependent ideas. As an indirect realist he takes himself to have three key commitments.

First, there is a mind-independent, or external, world filled with mind-independent objects that

² See, for example, discussion in Robinson (1994). The difference between these sometimes seems purely terminological. On some occasions, however, the direct objects of awareness may be given different characterizations. For example, some theorists argue that sense-data are or may be mind-independent, while some offer different accounts of their relation to external objects. It is not clear if we should think of these as different characterizations of the same thing or as different theories with different entities that are simply given the same label. Undoubtedly, certain expressions carry with them certain baggage due to the theories in which they feature. Using the expression “sense-data” brings with quite a bit of baggage and may invite more misunderstanding than alternatives. Thus, opting for alternatives can be useful.

have their properties and exist independently of being perceived. Second, perceiving an external object involves causally interacting with it. Ideas result from such interaction. Third, what we are directly aware of in perceptual experience are ideas. We are indirectly aware of external objects via being directly aware of the ideas caused by these objects.

Sollberger describes perceptual experience in the following way:

Suppose you are looking at a fruit bowl that contains green apples and yellow bananas. According to indirect realism, in having such an experience the subject is somehow most directly aware of a range of sensory or phenomenal qualities, such as greenness or yellowness. These sensory properties, which have been caused by the fruits via a complicated causal chain, are mind-dependent: they are such that their existence depends upon their being experienced. (p. 815)

On his view, such phenomenal qualities belong to ideas. How things appear to a person is determined by internal ideas. We are to understand our perceptual experience as being caused by external objects, but to involve the awareness of ideas that have sensory or phenomenal qualities. Furthermore, in virtue of the relationship between them, ideas represent external objects.

The subjective sensory experience thus constitutes a representation of the external material world and this is the reason why this view is also called “representative realism.” (p. 816)

I will stick to the label “indirect realism” when discussing Sollberger’s view. And in what follows below, when I consider indirect realism I have in mind the sort of view he defends. His decision to call the objects we are directly aware of in perceptual experience “ideas” as opposed to “sense-data” or something else appears to simply be a matter of preference. It is clear, though,

what his ideas are: they are inner, or mental, objects that represent the external world—that is, they are mental representations.

2.2. Critical Realism

According to critical realism, perceptual experience involves two components. There is a sensory component, which is a conscious internal state that possesses phenomenal qualities. There is also a conceptual component involving the exercise of classificatory concepts. We perceive external objects when we have experiences caused by those objects. An experience includes an inner state that is distinct from the object that caused it. What we are aware of in experience, or what appears to us, are the qualities of this state. That is to say that the phenomenal character of experience is determined by this inner state. At the same time, we directly perceive external objects thanks to the concepts deployed referring to these objects.

On such a view, when I sit at my table looking at my computer I see my computer, which is a mind-independent object. The sensory qualities that appear to me and which make up my experience—the colours and shapes and sounds—are mind-dependent. But while the phenomenal qualities I am aware of belong to an inner state, I also refer directly to the computer with the concept(s) exercised in perception.

Coates (2007) defends a critical realist view. On his view, which he characterizes as a type of causal theory of perception,

...in perceiving a physical object, the subject has a perceptual experience. This experience is an *inner* state that is, in important ways, distinct from, and caused by, the particular object perceived; yet in having an experience, we perceive objects directly. (p. 4)

It is the two components of the critical realist picture that allows Coates to both say that there are inner states “to which phenomenal qualities belong” that are distinct from physical objects, and that we perceive physical objects directly. He describes the conceptual component as a “perceptual taking” that refers to the objects that cause experiences. So while what appears to us in experience are mind-dependent phenomenal qualities, there is an important direct connection to external objects.

When I look in normal circumstances at a red apple, I am nonconceptually aware of a red phenomenal quality, which belongs to my inner phenomenal state. But I do not usually attend to my inner state, and it is not what I see. I see the apple, I classify it directly as such. (p. 10)

Although Coates does say that a perceptual experience is an inner state—see the quote above—he is careful to note that “the overall perceptual experience should not be identified with a purely phenomenal state” (p. 4). A perceptual experience contains a phenomenal component and a conceptual component. The phenomenal or sensory component is the inner state, which is a phenomenal state. It is due to this phenomenal state that the subject is aware of phenomenal qualities, which belong to this inner state. But in having the inner state there is also a conceptual component. Perceptual experience involves the exercise of concepts, in particular classificatory concepts. In perceptual experience, while having the inner state, there is also some conceptualization taking place.³

³ There is certainly no problem with an individual being in two states at once. The issue here is simply understanding what is going on when an individual is in a perceptual state. By saying that an experience is an inner state and also that the phenomenal component is an inner state, Coates may leave the reader wondering just where the conceptual component fits in. Is an experience made up of two states or is it one state with two components? Though he at one point describes the conceptual component as the “exercise of conceptual states” (p. 13), he typically talks about

Critical realist (or two-component, two-aspect, or two-factor) views have been defended by both Roy Wood Sellars (1939) and Wilfrid Sellars (1956), as well as more recently (Coates, 2007; Hatfield, 2016). Like indirect realism, critical realism takes perceptual experience to involve the awareness of internal phenomenal qualities. However, in contrast with indirect realism—where phenomenal qualities are taken to belong to mind-dependent objects we are aware of in experience—the inner sensory state is not understood as an object.⁴

2.3. Intentionalism

According to intentionalism, perceptual experience is a matter of mentally representing the external world. We perceive mind-independent objects when we represent them as being a certain way. Intentionalists argue that the phenomenal qualities we are aware of in perceptual experience are properties of mind-independent objects; they do not belong to some inner state or object, nor are we aware of something mind-dependent. However, the phenomenal character of an experience is due to the representational or intentional properties of the experience.

On such a view, when I sit at my table looking at my computer I see my computer. I represent the mind-independent object as having certain features, or as being a certain way. And

exercising concepts in experience. But if experience is one state with two components, we may wonder how we are to understand the conceptual component given that he says there is an inner state that is the phenomenal state (rather than simply saying there is a phenomenal component). I take it we should understand experience as including a phenomenal state and a conceptual component, as I've described here. Whether the conceptual component involves being in a conceptual state may not be important. We can view perception as a process that involves the exercise of concepts—the perceptual taking—and also results in an inner phenomenal state.

⁴ Adverbial views (see, e.g., Chisholm, 1957) also take the character of experience to be determined by qualities of an internal sensory state, while denying that perception involves a relation to some inner object(s). This renders such views similar to critical realism, though lacking the conceptual component. I will return to the difference between these approaches and indirect realism below.

the sensory qualities that appear to me and which make up my experience—the colours and shapes—belong to the computer.

Intentionalism has become fairly popular as of late, and a variety of intentionalist approaches have been recently defended (see, among others, Harman, 1990; Tye, 1992, 1995; Byrne, 2001). Because perception is understood in terms of representation—the view is sometimes even described simply as committed to the idea that experiences are representations—intentionalism is sometimes called “representationalism” (for example, in Crane, 2006; Genone, 2016; Berger and Nanay, 2016). However, I will stick to “intentionalism,” in no small part because there are other views that understand perceptual experience as involving representation or having representational properties.

In characterizing perception, Harman draws attention to what he calls “the intentionality of experience”:

Our experience of the world has content—that is, it represents things as being a certain way. In particular, perceptual experience represents a perceiver as in a particular environment, for example, as facing a tree with brown bark and green leaves fluttering in a slight breeze. (p. 34)

Perceptual experiences are construed in representational terms, as having intentional content, and are understood in terms of such content. The content of an experience is the way it represents objects as being.⁵

Intentionalists stress that what we are aware of in experience are not mental objects or qualities, but rather properties of external objects. For example, Harman considers someone, Eloise, seeing a tree:

⁵ This is also sometimes put in terms of “propositional content” (see, e.g., Byrne, 2001), but I will stick to “intentional.”

...what Eloise sees before her is a tree... That is to say, the content of her visual experience is that she is presented with a tree, not with an idea of a tree. Perhaps, Eloise's visual experience involves some sort of mental picture of the environment. It does not follow that she is aware of a mental picture. If there is a mental picture, it may be that what she is aware of is whatever is represented by that mental picture; but then that mental picture represents something in the world, not something in the mind. (p. 36)

The properties that appear to Eloise—for example, the colours of the tree—are properties of the tree, as it is represented to her. They are not the properties of some mental object. As Harman urges, we should not conflate “the properties of a represented object and the properties of a representation of that object” (p. 35). Similarly, Tye describes his standing on the beach looking out at the ocean as follows:

I found myself transfixed by the intense blue of the Pacific Ocean. Was I not here delighting in the phenomenal aspects of my visual experience?
 ... It seems to me that what I found so pleasing in the above instance, what I was focusing on, as it were, were a certain shade and intensity of the colour blue. I experienced blue as a property of the ocean not as a property of my experience. My experience itself certainly wasn't blue. Rather it was an experience that represented the ocean as blue. What I was really delighting in, then, were specific aspects of the content of my experience. It was the content, not anything else, that was immediately accessible to my consciousness and that had aspects I found so pleasing. (1992, p. 160)

Importantly, according to intentionalism, the representational or intentional properties of an experience completely determine its phenomenal features. A perceptual experience is a certain way, or has the phenomenal character it does, because of how the individual is representing the world.

2.4. Naïve Realism

According to naïve realism, perceptual experience involves the direct awareness of mind-independent objects. When we perceive external objects they are immediately present to the mind. Mind-independent objects and their properties are constituents of our perceptual experiences, and experience relates us to these objects. Thus, the phenomenal character of experience is due to mind-independent objects.

On such a view, when I sit at my table looking at my computer I see my computer, which is an external object. My experience is of my computer—I am aware of the object directly—and the sensory qualities that appear to me are properties of the mind-independent object.

Martin (2001, 2002, 2004, 2006) defends a naïve realist theory. On his view, “experience is a relation between the subject and some object of awareness” (2004, p. 42). However, in contrast to indirect realist theories which take perception to involve the awareness of something mind-dependent, the objects of awareness are in this case mind-independent objects. Like intentionalists, Martin thinks experiences are of mind-independent objects, however he denies that this is in virtue of the representational properties of experiences. Rather, “one’s experience relates one to the mind-independent world, and yet does so in a non-representational manner” (2002, p. 378). The objects present to the mind in experience “must actually exist and genuinely be present to the mind” (p. 393). Mind-independent objects and their properties are therefore

constituents of experiences and determine the phenomenal character of our perceptual experiences.

Naïve realism is now quite popular amongst philosophers, and a variety of naïve realist accounts have recently been offered in addition to Martin's (Campbell, 2002; Brewer, 2007, 2011; Fish, 2009). It is sometimes claimed that naïve realism is, or captures, our common sense view of perception. For instance, Brewer takes the claim that physical objects are the direct objects of perception to be an expression of the "commonsense idea that physical objects are the objects which are *presented* to us in perception, whose nature is made evident to us by the subjective character of our perceptual experience" (2007, p. 88). Crane, meanwhile, takes the idea that "perceptual experience is a genuine relation to its objects" to be "one of the most obvious or commonsense features of perception" (2006, p. 140-1). He also claims that "it is part of our commonsense conception of perception that the objects and properties we experience when we do perceive *are* the objects and properties out there in the mind-independent world" (2005, p. 246). I am not sure if it is right to view naïve realism as our common sense view of perception.⁶ I am not even sure there is anything like a common sense view about the nature of perception, much less a single, widely shared one. But, even if there was, its significance would be questionable insofar as we are in search of a correct theory about the nature of perception.

⁶ Genone's (2016) take here is more subtle. He describes direct realist views as ones on which, in experience, we are directly aware of mind-independent objects. Though he notes common sense could be mistaken, he argues that "a primary appeal of direct realism is common sense" (p. 4). In his taxonomy, naïve realism is a particular kind of direct realist approach, one with certain additional commitments. Here, Genone says "there is nothing commonsensical or naïve about naïve realism" (p. 2). Martin, for his part, is somewhat unclear. He notes that the person who defends the sort of theory he offers "claims to be doing justice to some common sense or naïve intuition about the kind of direct access to the world that perceptual experience can provide for us" (2002, p. 398). At the same time, he notes: "One might be sceptical whether it could really be part of any common sense view that objects were or were not constituents of our experiences of them" (p. 398). It is unclear, though, whether such reservations are his own, or those of a possible objector.

It may, however, be true that naïve realism captures how our perceptual experience seems to be. Nudds, for example, considers the experience of a vase:

Suppose that you are looking at a vase of flowers on the table in front of you. You can visually attend to the vase and to the flowers, noticing their different features: their colour, their shape and the way they are arranged. In attending to the vase, the flowers and their features, you are attending to mind-independent objects and features. ... [And] your experience seems introspectively to involve those mind-independent objects and features. (2009, p. 334)

When we have experience it seems like things are simply present to us—for example, a vase, or flowers, or a computer. From the point of view of a person having experience, it may seem like physical objects are immediately present, and that what is going on in perception is that we are in direct contact with the world and its objects.

Naïve realism is sometimes called “direct realism” or “relationalism.” However, some people distinguish between direct realist views and naïve realism (see, e.g., Genone, 2016). Meanwhile, though drawing attention to the relational component of naïve realism may help in contrasting it with intentionalism, indirect realism can also be seen as a relational view. Therefore, I will avoid using these other labels and stick with “naïve realism.”

3. Some Motivating Considerations

The considerations motivating the above theories are varied, and an exhaustive examination would not only be unwieldy but of questionable value. However, reviewing some of the main arguments that feature in the discussions of these theories can help deepen our understanding of the views and reveal the crucial differences between them.

3.1. The Scientific Causal Picture of Perception

There is a long scientific tradition that treats perception as comprising a series of stages, which in outline goes back to Descartes's work in *The Optics*. The idea on this view is that perception is a causal process, connecting the physical object in the environment to neural events in the perceiver's brain, and ultimately to an *inner* state of consciousness. In veridical perception, the subject is in a state of conscious awareness that is distinct from the perceived object. (Coates, 2007, p. 52-3)

As the scientific study of perception has advanced, accounts of the processes leading up to perceptual experience have been expanded and refined. Our understanding of the stages in perceptual processing and the brain structures in which they take place has deepened. But the general picture that emerges from empirical investigations of perception remains the same: causal interactions between our sense organs and the external world yield sensory input, which initiate a series of stages of processing in the brain, eventually resulting in perceptual experience.

We can now fill in this picture with endless details. Here are just a few, focusing on vision. Light is reflected off of or emitted by external objects and arrives at the retina. The retina has different types of cells, which respond differently, depending on the amount of light that hits them and the light's wavelength. The output of these rods and cones is processed by other cells in the retina, including retinal ganglion cells. Retinal ganglion cells then transmit a signal to the brain, their axons forming the optic nerve. The optic nerve carries the signal that was initiated by light to various parts of the brain, including the lateral geniculate nucleus, which then transmits the signal to primary visual cortex in the occipital lobe. In primary visual cortex there is a retinotopic map, meaning that the pattern of activations from the retina is preserved. As

information is transmitted to successive stages of visual processing in different areas of the brain, the receptive fields of neurons grow—a given neuron will respond to a larger portion of the visual field—and different kinds of information is processed. Neurons in later visual areas respond to more and more complex stimuli—for instance, to a complete object rather than to a line of a certain orientation. The brain’s processing of the visual input culminates in a visual experience. The particular brain areas responsible for our having different sorts of visual experiences—for example, of a face or a car—can even be observed in imaging studies.⁷

The indirect realist argues that this scientific causal picture supports their view of perception. Perception is a causal process involving many stages, taking place over a short period of time, that starts with external objects and results in our having experience. Increasingly complex stages of processing take us from a two-dimensional retinal image to a reconstruction of the external objects that caused the sensory input. The resulting perceptual experience supervenes on, or is possibly identical to, the brain state that represents the causes of the sensory input. While our perceptual experience is caused by external objects, it involves inner states that represent those objects and are distinct from them. My experience of a gray table has the properties it does because of the brain activity at the terminus of the string of processing. We can think of this brain activity as being, or as corresponding to, a mental object; and the properties I am aware of in experience—the colour and shape of the table—are the properties of that object. If the brain activity at the end of the processing was different and the mental object had different properties, then my experience would be different.

The critical realist also finds support for their view in the scientific causal picture. But rather than think of the result of perceptual processing in terms of a mental object, the critical

⁷ See Brown (2008, p. 50-3) for a similar discussion. He presents significantly more detail and focuses mostly on retinal processing.

realist argues that we should understand the result as mental state with certain properties. The inner state is distinct from any mind-independent objects that cause it. And it is the properties of this inner state that determine the character of our experience, for if the inner state differs—if the brain activity at the end of the chain of processing is relevantly different—then so will our experience.

3.2. The Argument from Transparency

The argument from transparency is offered in response to the notion that perception involves the awareness of some inner objects, states, or properties. Intentionalists and naïve realists alike argue that if we introspect on our perceptual experience we do not find any mind-dependent entities or inner states. Rather, introspection reveals only mind-independent objects like computers and cars.

The quotes from Harman and Tye above, in 2.3, appear in the context of the argument from transparency.⁸ Both Harman and Tye focus on what we are presented with in experience. For Tye, it is the blue of the ocean that he experiences. For Harman, Eloise’s experience presents the tree to her in a certain way—for example, as viewed from a certain location. Both deny that what is presented to us in experience is something mind-dependent.

We can see that this is the case, they argue, by introspecting on our experience. If we attend to our perceptual experiences, we will not find any inner objects or states. Our experience is transparent, or “diaphanous,” to the external world. Thus, in considering his sensation of blue, Tye notes:

⁸ The quote from Nudds in 2.4 expresses the same idea. However, he does not focus (in 2009) on articulating and defending the argument from transparency.

When one tries to focus on it in introspection one cannot help but see right through it so that what one actually ends up attending to is the real color blue. (p. 160)

The idea here is that if our experience involved the awareness of something mind-dependent we should be able to attend to the mind-dependent objects, states, or properties by turning our attention to our experience itself. However, as Harman argues, Eloise,

...is not aware of those intrinsic features of her experience by virtue of which it has that content. Indeed, I believe that she has no access at all to the intrinsic features of her mental representation that make it a mental representation of seeing a tree...

...I want to say that she is not aware of, as it were, the mental paint by virtue of which her experience is an experience of seeing a tree. She is aware only of the intentional or relational features of her experience, not of its intrinsic nonintentional features.

Some sense datum theorists will object that Eloise is indeed aware of the relevant mental paint when she is aware of an arrangement of color, because these sense datum theorists assert that the color she is aware of is inner and mental and not a property of external objects. But, this sense datum claim is counter to ordinary visual experience. When Eloise sees a tree before her, the colors she experiences are all experienced as features of the tree and its surroundings. None of them are experienced as intrinsic features of her experience. Nor does she experience any features of anything as intrinsic features of her experience. And that is true of you too. ...Look at a tree and try to turn your attention to intrinsic features of your visual experience. I predict you

will find that the only features there to turn your attention to will be features of the presented tree... (p. 39)

If we turn our attention inward to our experience, Harman argues, the qualitative features of our experience do not change. What is present to the mind in perception are the properties of mind-independent objects like computers and trees, as the intentionalist claims. The sensory qualities we are aware of are properties of such objects. And turning our attention to our experience does not make us aware of some inner phenomenal properties—phenomenally our experience does not change, and the features we are aware of continue to be those of external objects.⁹

Martin recognizes that the argument from transparency has been offered in defense of intentionalism, however he thinks the argument also supports naïve realism. If the choice is between intentionalism and views on which experience involves the awareness of something mind-dependent, the argument from transparency supports the former. But, he notes, naïve realism affords us another option. In fact, he argues that transparency offers even better support for naïve realism.

His presentation of the argument parallels what we see in Harman and Tye. He considers the case of his looking at the lavender bush on his street:

When I stare at the straggling lavender bush at the end of my street, I can attend to the variegated colours and shapes of leaves and branches. ...When my attention is directed out at the world, the lavender bush and its features occupy center stage. It is also notable that when my attention is turned inwards instead to my experience, the bush is not

⁹ Things may change cognitively when we attend to our experience itself, but the argument here is that things remain constant phenomenally, i.e. when it comes to the phenomenal properties present—for example, the colours and shapes.

replaced by some other entity belonging to the inner realm of the mind in contrast to the dilapidated street in which I live. I attend to what it is like for me to inspect the lavender bush through perceptually attending to the bush itself while at the same time reflecting on what I am doing. So it does not seem to me as if there is any object apart from the bush for me to be attending to or reflecting on while doing this. (2002, p. 380-1)

While turning his attention inward allows him to reflect on what he is doing, things do not change perceptually when he does this. And while the lavender bush and its properties are presented to him in experience, there are no mental objects, and no inner state, when he searches his mind.

The naïve realist accounts for the transparency of experience differently than the intentionalist. It is not the experience's intentional content that explains transparency, but that the objects we perceive are actually present in experience. One may think the naïve realist thus offers a better account of our experience than the intentionalist. This is, in fact, what Martin claims: "I suggest that we should think of Naïve Realism as the best articulation of how our experiences strike us as being to introspective reflection on them" (2004, p. 42).¹⁰

The proponent of the argument from transparency thinks that indirect realists and critical realists are wrong to conclude on the basis of the scientific causal picture of perception that perception involves the awareness of inner objects or states. Such an interpretation of the

¹⁰ It is worth noting that Martin (2002) offers an argument from transparency against intentionalism, focusing on sensory imagination instead of perceptual experience. He argues that intentionalism faces difficulties when confronted with introspective evidence of imagining, in much the same way that certain theories allegedly face difficulties with introspective evidence of experience. I do not think that in the preceding quote he is referring both to experience and imagining with "experiences." I may, of course, be wrong here. However, Martin suggests in a number of places that naïve realism best captures how our perceptual experiences seem to us (see, for example, 2004, p. 51).

scientific picture, it is implied, is incorrect. Since the argument from transparency is taken to show that there are no inner objects or states, concluding otherwise on the basis of scientific work on perception is mistaken. After all, intentionalists and naïve realists are not keen to dispute the sorts of neuroscientific details mentioned above.

There are a number of possible responses to the argument from transparency. Perhaps the most obvious point to make is that it is not clear why we should expect any phenomenal change when we introspect on an experience. The indirect realist, for example, argues that in experience we are aware of mind-dependent objects and properties, not that we are aware of mind-dependent objects and properties as mind-dependent objects and properties. The view is that the mind-dependent objects and properties determine the phenomenal character of our experience. The indirect realist need not ever claim that there will be a qualitative change when we introspect on an experience, nor that we can attend to the mind-dependent properties so as to become aware of their mind-dependence.

It is also not clear what we should take to follow from how things seem to us. The indirect realist or critical realist may be willing to concede that in experience it seems to us like external objects like computers and trees are present, or that experience involves these objects, but think that this tells us little about the nature of perception. Unless what we want from our theory of perception is simply to describe how things seem to us in experience, it is not clear what truths about the nature of experience and our relationship to mind-independent objects follow from how things seem to us.¹¹

¹¹ Hatfield (2007) makes this point while defending a particular view about the nature of colour. He notes that some philosophers:

...expect our experience to be *metaphysically transparent*: if we attentively inspect our sensory experience, it will reveal its real nature to us. However, there is no reason to suppose that the senses have evolved in order to reveal the principles of their own operation, or to

In particular, we may hesitate to draw conclusions about how things actually are based on how they seem if we can explain why it is that they seem the way they do. Such an explanation is far from unimaginable. For example, Metzinger (2009) defends an indirect realist view of perception while offering an explanation of why things may seem to us to be as the naïve realist describes. Contemporary accounts of motion, matter, and mind are often in tension with how things may seem to us from certain perspectives, yet we have proven adept at resolving such tension.

3.3. The Traditional Arguments from Illusion and Hallucination

The arguments from illusion and hallucination have been taken to demonstrate that in perceptual experience we cannot be immediately aware of ordinary external objects.

Traditionally, the arguments have been advanced in support of sense-data. Though the arguments and their conclusions would likely have been accepted by many modern philosophers, they only appear in earnest in the early twentieth century (Robinson, 1994). The argument from illusion grows out of the simple idea that objects sometimes look, feel, smell, taste, or sound to us differently from the way they actually are. The argument from hallucination raises the stakes further, considering cases where we have an experience in the absence of an object.

The argument from illusion can be constructed around any case where it seems to one that an object has a certain property which it lacks. The straight stick that when partly submerged in water appears bent has been generous in its service. Other familiar examples are objects appearing to be of different shapes when viewed from a distance or at an angle, such as a cylindrical tower appearing rectangular or a round coin elliptical. Objects may also look to be of

suppose that the true physical descriptions of the object properties that cause (and are signified by) our sensory representations should be transparently available in consciousness. (p. 149)

a different colour when viewed from a certain perspective, such as mountains appearing purple from a distance. I will not devote time here to considering which of these classic examples are genuine instances of illusion. If one is not happy with the cited case, they should feel free to insert their favourite illusion into the argument. In the interest of not straying from tradition I will focus on the case of a white wall in red light such that it looks as though it is painted red.

Here is the argument from illusion (adapted from Crane and French, 2015):

- (i) In an illusory experience it seems to one that something has a property, F , which the mind-independent object perceived lacks. Though the wall seems red to me, it is actually white, not red.
- (ii) When it seems to one that something has a property, F , then there is something of which one is aware that has that property. Traditionally, sense-data are invoked here, as the objects we are aware of. Thus, what I am aware of is a red sense-datum.
- (iii) Since the mind-independent object in question does not have the property, F , then in illusory experiences one is not aware of the mind-independent object.
- (iv) The same account of experience must apply to both veridical and illusory experiences.
- (v) Therefore, one is never perceptually aware of mind-independent objects. What one is aware of in experience are sense-data. (Or, if we like, we are directly aware of sense-data and only indirectly aware of mind-independent objects.)

The key premises in the argument are (ii) and (iv). Premise (ii) has been labelled the “phenomenal principle” by Robinson (1994, p. 32).¹² The motivation here is explanatory: since what appears before me is red, we cannot explain this by saying that I am aware of the white wall; rather, it seems like I am aware of something else, something red. Premise (iii) is taken to

¹² Smith (2002, p. 25) calls this the “sense-datum inference.”

simply follow from the preceding considerations. The very reason we invoke sense-data is because claiming that I am aware of the mind-independent object would not explain my experience. Instead, we say that I am aware of sense-data. Therefore, I am not aware of the white wall.¹³

Premise (iv) is the spreading or generalizing step that takes us from the case of illusion under consideration to all cases of perceptual experience. The basic idea is that sense-data are invoked in order to explain the experience in question—for example, my experience of a red wall when I am before a wall that is actually white. It is then claimed that we should give the same account for all perceptual experiences, whether veridical or illusory. This claim is often supported in one of two related ways. First, since these experiences are subjectively indistinguishable we should understand them as having the same nature. Second, veridical and illusory experiences form a continuous stream, and such continuity would be broken if different experiences had different natures, or if experience alternately involved the awareness of mind-dependent and mind-independent objects.

The argument from hallucination has the same structure as the argument from illusion, but is constructed around a case of hallucination. For example, suppose I have the experience of seeing a sea turtle on the floor in front of me when there is nothing there:

¹³ Crane and French (2015) argue that an additional premise is needed to move from (ii) to (iii). Typically, (iii) is taken to follow due to Leibniz's law (see also Robinson, 1994; Smith, 2002). The wall is white, not red, but what I am aware of is red, not white. Since what I am aware of and the wall have different properties, they cannot be identical, therefore I cannot be aware of the wall. However, to exclude the possibility that I am aware of a sense-datum and the wall a further premise is needed: that if I am aware of a red sense-datum I am not also aware of the white wall. I do not find Crane and French's point here compelling, since the claim that one is aware of both things is difficult to understand given the role of direct objects of awareness, i.e. as determining the character of experience. And there is simply no reason to claim that in the illusory case one is also aware of the mind-independent object that does not have the relevant property. However, since I am not defending the traditional argument from illusion I will not explore this further.

- (i) In a hallucinatory experience it seems to one that there is some object, *O*, when there is no such mind-independent object. While it seems to me like there is a sea turtle in front of me, there is actually nothing there.
- (ii) Since there is no mind-independent object present, then in hallucinatory experiences one is not aware of a mind-independent object.
- (iii) When it seems to one that there is an object present, then there is something one is aware of. Traditionally, sense-data are invoked here, as the objects we are aware of. (Alternatively, we may say that I am aware of a mental image of a sea turtle.)
- (iv) The same account of experience must apply to both veridical and hallucinatory experiences.
- (v) Therefore, one is never perceptually aware of mind-independent objects. What one is aware of in experience are sense-data. (Or, if we like, we are directly aware of sense-data and at most indirectly aware of mind-independent objects.)

Premises (ii) and (iii) above can be flipped, but the way I have presented them here strikes me as more natural. Premise (iii) is the phenomenal principle and premise (iv) is the spreading premise. Once again, the idea is that in order to explain cases of hallucination we say we are aware of sense-data, then it is claimed that we should explain all perceptual experiences in the same way, thus we conclude that we are only ever aware of mind-dependent objects and never mind-independent objects (except possibly indirectly so).

While these arguments were first advanced in support of sense-data theories, we do not have to appeal to sense-data in the course of either argument. Instead, we can point to mind-dependent ideas or inner states that possess phenomenal qualities. The allegation is that we cannot explain what appears to one in these cases—i.e. the phenomenal character of these experiences—with mind-independent objects. The challenge of explaining the phenomenal

character of illusory and hallucinatory experiences can be met equally well by claiming we are aware of ideas or the properties of inner states. It is worth noting, too, that cases of illusion and hallucination need not involve our being deceived by the experiences. The central issue is simply about explaining what the experiences are like, or the phenomenal properties that appear to us.

Both the phenomenal principle and the spreading step have faced significant pushback. I will not try to adjudicate the disputes about these premises or thoroughly evaluate the traditional arguments from illusion and hallucination. However, briefly examining the responses to these two key premises can help us gain a fuller understanding of the different views of perception under discussion.

The phenomenal principle, it has been argued, involves a simple fallacy. It is simply wrong to think, the reply goes, that when I experience the red wall there need be any object I am aware of to which the property red belongs. Harman argues that “the argument fallaciously infers” that there must be some mental red I am aware of (1990, p. 35). The intentionalist argues that we do not need to posit some mind-dependent object or inner state we are aware of, as the indirect realist or critical realist would, but can instead explain the experience in terms of how I am representing the wall. That avenue may appear promising. The question is whether we can adequately explain illusions and hallucinations without invoking some internal objects or states to which phenomenal qualities belong.

The spreading step, meanwhile, has simply been rejected by naïve realists. Naïve realists deny that we need to give the same account for veridical experiences and hallucinations because they deny that these states are of the same fundamental kind. Instead, they embrace disjunctivism. So while it may be possible that it is indistinguishable to me whether I am hallucinating a sea turtle or actually seeing one, it does not follow that we need to account for

both experiences in the same way. Even if we are not directly aware of mind-independent objects in cases of hallucination—since by definition no such objects are present—it does not follow that we are not directly aware of mind-independent objects in veridical cases. By denying that hallucination and veridical perception are of the same kind and should be understood in the same way, the naïve realist can maintain that in perception we are directly aware of external objects.¹⁴

We have already seen how our different theories of perception offer differing accounts of our perceiving mind-independent objects. We can now see how they also offer different accounts of illusions and hallucinations.

An indirect realist like Sollberger argues that perceptual experience involves the awareness of mind-dependent ideas. In cases of veridical perception and illusions, these ideas are caused by mind-independent objects and events.¹⁵ When we experience hallucinations we are aware of mind-dependent ideas that are not caused by the relevant external objects but have some other origin.

¹⁴ As I note below, naïve realists may also reject the spreading step when it comes to the argument from illusion. Some naïve realists, however, may accept that illusions and veridical perceptions are of the same fundamental kind, but reject the phenomenal principle in the case of illusions, arguing instead that in such cases we are only aware of mind-independent objects. As sense-data theories have fallen out of fashion and the argument from illusion has seen its reputation sag in recent years, attention has shifted to hallucinations. Hallucinations have become more of a focus in the philosophical literature, in particular in work debating naïve realism. Here, the discussion often revolves around the possibility of hallucinations that are subjectively indistinguishable from experiences of objects, and that by definition do not involve the presence of a mind-independent object (see, e.g., Crane, 2005).

¹⁵ There is the difficult question of when an experience is veridical, or where to draw the line between veridical and non-veridical experiences. We may think a similar question arises regarding illusions and hallucinations, as some experiences typically thought of as illusory may seem to have a hint of hallucination. Precisely when an experience is veridical, illusory, or hallucinatory is not too important here. The three-part distinction is widely accepted, and for our present purposes we can simply discuss the three sorts of cases without being concerned with exactly how to characterize them or draw the lines between them. These questions are equally important for all theories of perception. I will focus more on this issue in the final chapter.

The critical realist strategy is similar. Veridical experiences and illusions involve having an experience caused by mind-independent objects, where the experience includes an inner state with phenomenal qualities. When we hallucinate we are also aware of the qualities of an inner state, but one that was not caused by external objects.

Intentionalists argue that we perceive mind-independent objects when we represent them as being a certain way. They claim that what we are aware of in cases of veridical perception are mind-independent objects and their properties. Despite this claim, intentionalists take the intentional properties of an experience to determine its phenomenal character. This may strike us as puzzling. If the phenomenal qualities we are aware of in perceptual experience are properties of mind-independent objects, in what sense do the intentional properties of the experience determine its phenomenal character? I will return to this question below. But, in any case, we can see how intentionalists will account for cases of hallucination. It is possible for us to be in an intentional state even if the object that state purports to be about does not exist. Similarly, it is possible to be in a state with the same intentional content whether or not the same object is before us. Thus, Harman says:

...what Eloise sees before her is a tree, whether or not it is a hallucination. That is to say, the content of her visual experience is that she is presented with a tree... (p. 36)

In claiming that it is the intentional properties of our experience that determines its phenomenal character, the intentionalist can say that a veridical perception and a hallucination are indistinguishable if they have the same intentional content. It is simply that in cases of hallucination we represent the world as being a way it is not. Meanwhile, the intentionalist claims that illusions arise when we represent mind-independent objects as being a way they are not. But it is fair to wonder if more needs to be said here. If the phenomenal qualities we are

aware of in perception are properties of mind-independent objects, then cases of illusion and hallucination would seem to present a puzzle. In such cases our experiences certainly have phenomenal qualities, but the relevant mind-independent objects or properties are not present.

Naïve realists take perceptual experience to involve the direct awareness of mind-independent objects. When we perceive external objects they are immediately present to the mind and they determine the phenomenal character of our experience. The same account cannot be given for hallucinations because there are no mind-independent objects present. This is why naïve realists embrace disjunctivism. Naïve realists must reject the spreading step and argue that veridical and hallucinatory experiences are fundamentally different kinds of mental states with different natures. The disjunctivist claims that when I have the perceptual experience of a sea turtle this may be a case of veridical perception or of hallucination. If it is the former, the naïve realist will explain the character of my experience in terms of the sea turtle present to my mind. Whether naïve realists can offer a compelling account of the case where I am hallucinating is an interesting question. As far as the naïve realist is concerned, though, the crucial point is that however one wants to explain my hallucination of the sea turtle, the same sort of account will not apply to the case where I actually see a sea turtle. The naïve realist avoids the conclusion of the argument from hallucination by accepting disjunctivism.

Many naïve realists will also embrace disjunctivism in response to the argument from illusion. This means grouping illusions and hallucinations on the other side of the disjunction as veridical perception. Martin adopts such a position (see, e.g., 2002, p. 393-9; 2004, p. 44). The appeal of this approach is easy to see. Since illusions are cases where a person has an experience of an object having some property the object actually lacks, it is hard to see how the naïve realist can account for the character of that experience by pointing to the object itself. Placing illusions

alongside hallucinations allows the naïve realist to avoid the conclusion of the argument from illusion by denying the spreading step. However, some naïve realists group illusions and cases of veridical perception together (see, e.g., Fish, 2009; Brewer, 2011; for a discussion of different varieties of disjunctivism, see Byrne and Logue, 2009). Those adopting this strategy must offer their own account of illusions. Whether naïve realists can offer a compelling explanation of illusory experiences that is consistent with the core commitments of their theory is something I will consider in a later chapter.

3.4. The Causal Argument for the Common Kind Assumption

The traditional arguments from illusion and hallucination take as a premise the claim that the same account should be offered for cases of veridical perception, illusion, and hallucination. We can understand the thought implicit here as being that all perceptual experiences are of the same fundamental kind. Because these different cases are of the same kind, they ought to be given the same account; if they were importantly different in their natures, the same account need not apply to all of them. The idea that veridical perception, illusion, and hallucination are fundamentally the same kind of mental event is called the “common kind assumption.” The common kind assumption is so closely linked to the spreading premise that we may think of the common kind assumption as being a premise in the traditional arguments from illusion and hallucination.

Naïve realists reject the common kind assumption. In doing so, they accept disjunctivism. As we have seen, this move is made in order to save naïve realism from the argument from hallucination. Martin explicitly acknowledges this: “The prime reason for endorsing

disjunctivism is to block the rejection of a view of perception I'll label *Naïve Realism*" (2004, p. 38). Disjunctivism, he says,

...seeks to resist the rejection of Naïve Realism, therefore, simply by denying the Common Kind Assumption. That is, we hold on to Naïve Realism by insisting that the fundamental kind of event that one's sensory experience which is a veridical perception of the table in front of one is a kind of event which just could not occur were one hallucinating.

(p. 43)

Rejecting the common kind assumption makes one a disjunctivist.¹⁶ Accepting the common kind assumption means supporting a non-disjunctive view of perception. Since the naïve realist has no choice but to reject the common kind assumption and be a disjunctivist, an argument in support of the common kind assumption is also an argument against naïve realism.

Sollberger (2008, 2012) presents a causal argument in support of the common kind assumption. It is an argument in support of the claim that the same sort of account must be given for hallucinations, illusions, and veridical perceptions.¹⁷

¹⁶ As I will argue below, rejecting the common kind assumption is required in order to preserve the sort of directness essential to naïve realism. For this reason, by "disjunctivism" I will mean a naïve realist disjunctivism. Conceivably, one could hold an alternative view about the nature of perception and be a disjunctivist. However, the motivation for doing so is unclear, unless one's view involves the same sort of directness as naïve realism—but, in this case, however the views may differ they would share important core commitments (i.e. those of naïve realism).

¹⁷ Similar arguments are found in Robinson (1994, p. 151) and Coates (2007, p. 53). While those who accept the common kind assumption take veridical perception, illusion, and hallucination to form a common kind, the causal argument focuses on veridical perception and hallucination. However, if the argument succeeds, then naïve realism is defeated, and it is hard to see how someone could concede that veridical and hallucinatory cases were of the same fundamental kind but deny that illusions were. Furthermore, since many naïve realists group illusions with veridical cases, they accept that the two are the same fundamental kind of mental state. The argument thus targets all naïve realist views, regardless of how illusions are classified, or regardless of the variety of disjunctivism adopted; and, we can take the conclusion to also include illusions.

The causal argument asks us to consider a case where a person has a veridical perceptual experience and a case where a neuroscientist stimulates that person's brain in such a way that they have a hallucination that is subjectively indistinguishable from the veridical experience. The idea is that in both cases the person's brain will be in the same state, that the relevant brain activity in both cases will be identical. It is typically claimed that the two experiences therefore share the same neural cause; and, a principle like "same proximate cause, same immediate effect" is appealed to by proponents of the argument. However, formulating the argument this way suggests a certain sort of relationship between the physical and the mental: a causal one. But rather than thinking that brain states cause corresponding mental states, we may think the relationship between the physical and the mental is of a different nature—e.g. that it involves supervenience or identity. Since it does not seem like the argument should depend on the relationship between the brain and mind being a causal one, a slightly different formulation is preferable.¹⁸ This way, the argument cannot be rejected if it turns out that a certain relationship between the physical and the mental holds.

I have adapted Sollberger's argument to avoid suggesting what may be a problematic view of the relationship between the brain and mind. I will talk about a brain state, w , which an individual is in while having some perceptual experience. This is a brain state which is caused by another brain state, v , which in turn can be downstream of different causes. I take the addition of v to be unproblematic. Putting things this way preserves a common neural cause for the veridical and hallucinatory experiences, which is significant given that discussion of the causal argument typically focuses on a principle like "same cause, same effect."

¹⁸ Coates, for example, puts things in terms of a brain state being sufficient for an experience, whether the brain state is caused by an external object or has some other cause.

Consider my seeing a sea turtle and a neuroscientist stimulating my brain in such a way that I have a hallucination that is indistinguishable from the veridical perception of a sea turtle.

(i) Appropriate causal stimulation of my brain is sufficient for a hallucination of a sea turtle, h , to occur, which is subjectively indistinguishable from the veridical perception of a sea turtle. In this case, the causal stimulation results in my brain being in state, v , which causes brain state w , which I am in while having hallucination h .

(ii) My hallucinatory experience narrowly supervenes on my brain state. The occurrence of h requires only that my brain is in state w .

(iii) When I veridically perceive a sea turtle there is a causal chain beginning with the object and involving a series of brain states, eventually leading to a brain state, v , which causes another brain state, w , which I am in while having the perceptual experience. My perception, p , involves the causal chain that begins with the sea turtle interacting with my sense-organs and ends with w .

(iv) Although h and p have different aetiologies, the causal chains leading to h and p partly overlap; namely, both causal chains end with v followed by w . Brain state v brings about w , and is therefore sufficient for h .

(v) Therefore, h must be co-present with p while I perceive the sea turtle.

(vi) Given that h grounds my seeming awareness of a sea turtle, or the appearance of a sea turtle in experience, it is sufficient for grounding the phenomenology of p .

(vii) Therefore, h pre-empts p from explaining my mental state of perceiving a sea turtle, rendering p explanatorily idle. Since p is explanatorily idle it can be discarded; it is not the external object—the sea turtle—that determines the phenomenal character of my experience.

(viii) The fundamental mental kind common to both my hallucination and perception of a sea turtle is h (i.e. the mental state which supervenes on my brain state w , which in turn can be brought about in a variety of ways).

The conclusion of this argument is that the common kind assumption is correct. Veridical perceptual experiences and hallucinations involve the same fundamental kind of mental state and should be given the same account because both are due to the same proximal brain state. My being in a certain brain state while hallucinating a sea turtle accounts for the phenomenal character of that experience; my experience has the phenomenal character it does simply in virtue of my being in that brain state. Therefore, when I am in the same brain state as a result of causally interacting with an actual sea turtle, my being in that brain state can, alone, account for the phenomenal character of my perceptual experience. Though the causal chains in the veridical and hallucinatory cases differ, they partly overlap at the crucial point: they both involve a brain state that causes the brain state that I am in while having the perceptual experience of seeing a sea turtle.¹⁹ In addition to the veridical perception and the hallucination involving the same brain state and being indistinguishable, both can play the same psychological role and lead to the same behaviors.

The causal argument is neutral about just how to “construe the metaphysical nature of the fundamental kind” (Sollberger, 2008, p. 6). What it shows is that our perceptual experience supervenes on an inner state that is distinct from the mind-independent objects that may cause it. But it leaves open whether we should understand such states in terms of intentional objects, or as

¹⁹ This allows us to say that the two experiences have the same neural cause. Without v we would not necessarily be able to say this. For instance, we can imagine a neuroscientist stimulating my brain such that I am immediately in state w . But, in this case, we may consider the stimulation (or the neuroscientist), and not w , as the cause of my experience.

having mind-dependent phenomenal properties, or as involving the awareness of some mental object(s), or in some other similar way.

Martin accepts that experiences, like other states or events in the world, are subject to the causal order. He also accepts the possibility of the sorts of hallucinations described in the causal argument. But while he is willing to concede that “what is present in hallucination is also present in perception,” he denies that “what is present in perception is there in hallucination” (2004, p. 58). In particular, he argues that both a veridical experience and the corresponding hallucination share the property of being indistinguishable from the veridical perception. However, only the veridical perception has the property of being a veridical perception. Moreover, the explanatory role of being a veridical perception is preserved, as it appears essential to characterizing experiences as being indistinguishable from the veridical perception.

According to Martin’s view, my veridical perceptual experience of the sea turtle and my hallucination of the sea turtle are instances of the same kind of mental event. But the same mental event can be of two different kinds. The veridical perception is, more fundamentally, of a second kind, a kind which excludes the hallucination.

Clearly for a veridical perception, being a veridical perception of a tree is a better candidate for being its fundamental or essential kind than being indiscriminable from being such a veridical perception. When we turn to the case of the hallucinatory experience, there are no other candidates for the kind of mental event it is—at least according to the disjunctivist—other than its being indiscriminable from the veridical perception; *faute de mieux* this is then the fundamental mental character of the event. (p. 72)

Whether we find Martin's response to the causal argument compelling may depend on whether we agree with his characterizations of perceptual experience and of hallucination. As is clear from the quote above, his argument depends on how the disjunctivist characterizes experience. For Martin, being indiscriminable from veridical perception is central:

If we are to settle the concerns here, we need to make clear what can be, and what need not be, in common among matching perceptions and hallucinations. Here the notion of sense experience in general, that of being indiscriminable from veridical perception, plays a central role. ... For, the concept of perceptual experience in general is that of situations indiscriminable from veridical perception... (p. 64-5)

Regarding hallucinations, Martin offers an epistemological account, arguing that when it seems to me like a sea turtle is present while I hallucinate, we should understand "seems" here simply as saying that I would be unable to tell the situation apart from one where I was seeing a sea turtle.

Sollberger, for his part, is unmoved by Martin's response to the causal argument. He argues that whatever separates p from h is explanatorily idle. The hallucination is not only indistinguishable from the veridical perception, but the two are apt to play the same role in our psychological lives. And the best explanation for the similarities between the two states is that they are of the same fundamental mental kind. A veridical perception and a hallucination differ, for Sollberger, but not in terms of their mental properties.²⁰

The basic idea behind the causal argument is simple enough. Some people think we can have the same perceptual experience whether an object is present or not, or whether or not we are

²⁰ See Nudds (2009) for a defense of Martin's approach in responding to the causal argument. For a discussion of various disjunctivist responses to the causal argument, including Martin's, see Robinson (2013).

veridically perceiving a mind-independent object. If the brain activity is the same—the same proximate brain state is present—then we should give the same account of the two indistinguishable experiences. If this is right, then the character of our experience cannot be due to the mind-independent object, as the naïve realist claims; and, perception cannot involve our being directly related to mind-independent objects such that they are constituents of our experiences. As we have seen already, indirect realists, critical realists, and intentionalists think the same sort of account should be given for the two experiences. That is, unlike the naïve realist, they are happy to accept the common kind assumption.

4. Indirectness and Mental Representation

The common kind assumption typically features in the following sort of reasoning. In the case of hallucination a person has the experience of an external object being present, but there is no external object, therefore we cannot explain this experience or its properties by invoking the relevant external object. In the case of illusion a person has the experience of an external object being present and having certain properties, but there is no external object with such properties, therefore we cannot explain this experience or its properties simply by invoking an external object that lacks the relevant properties. Since veridical perceptual experiences, hallucinations, and illusions are fundamentally the same kind of mental event—i.e. perceptual experiences—having the same nature, we cannot explain the nature of perceptual experience by appeal to external objects. Why think perceptual experiences form such a common kind? Here, we may point to their being subjectively indistinguishable, or forming a continuous stream, or playing the

same role in our psychological lives.²¹ Or, following the causal argument, we may argue that the same brain state can underlie identical experiences, whether a veridical perception or hallucination. That is, we may say that we can be in the same brain state, w , whether an object is present or not, and this brain state underlies our perceptual experience in the sense that it is responsible for the phenomenal qualities that appear, as well as the other properties of the perceptual state. In a nutshell, since we can have the same experience without the relevant mind-independent object being present, then that object cannot be essential to the experience, cannot be part of its nature or what gives the experience its phenomenal qualities. Causal interaction with a mind-independent object may be a way to end up with a perceptual experience, but the object is distinct from the experience.

Accepting the common kind assumption has clear consequences for the role of mind-independent objects in perception and for how we are related to mind-independent objects. According to naïve realism, mind-independent objects are directly present to the mind, are constituents of our experiences. On such a view, we are very closely related to mind-independent objects in perception. We are, perhaps, as closely related to such objects as is possible. (I cannot fathom a closer, or more direct, relationship.) We do not perceive such objects by being in a mental state that possesses certain phenomenal or representational properties. Rather, mind-independent objects and their properties are simply present to our minds. Preserving this direct relationship requires rejecting the common kind assumption.

²¹ That an illusory experience can play the same psychological role as a veridical experience does not mean that it will, or that it will continue to into the future. For instance, one may later regard the experience as a bad case, to be ignored. And though we may be fairly adept at determining which of our experiences belong in the bad category, it is certainly possible for us to regard a veridical perception as a bad case.

In contrast to naïve realism, on views that accept the common kind assumption there is a sort of distance between ourselves and mind-independent objects. Non-disjunctive theories do not grant mind-independent objects as direct a role in perception. Our relationship to these objects in perception is not simply a matter of their appearing to us or being immediately present to the mind. Instead, we are causally and representationally related to such objects. But such relations lend mind-independent objects a certain remoteness—they are at the opposite end of the causal sequence, they are the thing being represented by a mental state. In saying there is a distance between ourselves and mind-independent objects, I mean simply that since the same sort of account will be given for veridical perceptual experiences and non, mind-independent objects cannot be as closely related to us and our minds—as fundamentally involved—in the veridical case as they are on naïve realism. Non-disjunctive theories take perception to involve inner states that are distinct from mind-independent objects such that those objects cannot be constituents of the mental states.

Fish (2004) argues that we should understand the direct/indirect distinction in terms of the distinction between disjunctive and non-disjunctive theories. The “underlying theoretical basis of the direct/indirect distinction,” he argues, is the following:

...what makes a theory indirect is the fact that the canonical characterizations of veridical perceptual experiences are such that experiences *of the very same kind* can be enjoyed even when the experience’s apparent object is absent. (p. 6)

The canonical characterization of an experience is the sort of theoretical analysis of an experience we have seen in section 2, what Fish calls “the proper philosophical characterization of the experience” (p. 5). We can understand it as our account of the nature of the experience.

Thus, if we consider a traditional sense-data theory:

...perceptual experiences are canonically characterized as sensings of certain collections of sense-data—when we perceive a red apple, for example, the perceptual state we are in would be canonically characterized as the sensing of red, apple-shaped sense-data. And as a hallucination of a red apple is explained as the sensing of red, apple-shaped sense-data *in the absence of* a real physical red apple, we can see that in both the veridical and non-veridical cases, the subject is held to have experiences of the same canonical kind. (p. 6)

In contrast, on a direct theory, even if a veridical perception of an apple and a hallucination of an apple may be indistinguishable, the former is the kind of experience which we can only have when an apple is present. Hence, “the direct theorist must therefore accept that the *canonical characterization* of subjectively indistinguishable experiences...might differ from one another” (p. 6). A theory is indirect if it accepts the common kind assumption and says we are in the same kind of state in veridical and non-veridical cases; a theory is direct if it rejects this, i.e. if it is a disjunctive theory.

By following Fish’s way of drawing the distinction we can highlight the crucial way that views accepting the common kind assumption differ from disjunctive views regarding the role external objects play in perception. Naïve realism takes us to be directly related to mind-independent objects in a manner that clearly contrasts with other theoretical approaches. Accepting the common kind assumption is incompatible with such directness. This is evident if we consider the canonical characterizations of veridical and hallucinatory experiences on these other theories.

The indirect realist takes perceptual experience to involve the awareness of mind-dependent objects. For Sollberger, the mind-dependent objects are ideas that possess phenomenal

properties. When I have a veridical perceptual experience I am aware of an idea possessing certain phenomenal qualities. If I have an indistinguishable hallucination I am aware of an idea possessing the same phenomenal qualities. The two cases differ in their causal origin, but they are nevertheless taken to be of the same in this respect: both involve the awareness of an idea, which determines their phenomenal character.

The critical realist takes perceptual experience to involve an inner state that is distinct from the external objects that may cause it. When I have a veridical perceptual experience I am in a state that possesses certain phenomenal qualities. If I have an indistinguishable hallucination I am in a state that possesses the same phenomenal qualities. Once again, the two cases differ in their causal origin, but they are of the same kind.²² In both cases the phenomenal character of the experience is determined by the inner state.

The intentionalist takes perceptual experience to involve mentally representing the world to be a certain way. When I have a veridical experience I am in a state with certain intentional properties, a state that represents the world as being the way it actually is. If I have an indistinguishable hallucination I am in a state with the same intentional content, however in this case the world is not at all as I represent it as being. In both the veridical and hallucinatory cases it is the intentional properties of my perceptual state that determine the phenomenal character of my experience. Things appear to me a certain way, to have certain properties, due to how I represent the world.

Though the naïve realist grants that we can have hallucinations that are indistinguishable from veridical perceptual experiences, the two cases are given radically different characterizations. In particular, only in veridical cases are mind-independent objects present to

²² See, for example, Coates (2007, p. 13).

the mind, and the phenomenal character of these—and only these—experiences is determined by mind-independent objects and their properties.

The sense in which non-disjunctive theories are indirect is brought out by how these theories understand perceptual experience, and how their accounts of the nature of experience differ from the naïve realist approach. However, some non-disjunctive theorists—notably intentionalists, and perhaps some critical realists—may object to their views being labelled as indirect. I think this is due to the fact that “direct” and “indirect” are used in a variety of different ways in the literature regarding perception. These labels may be used to denote different epistemological, metaphysical, or even conceptual relations to mind-independent objects, to characterize different sorts of acquaintance relations, and also to indicate the sorts of objects we can be said to perceive.

Views are often labelled “direct” to contrast them with indirect realist sense-data theories that claim we directly perceive sense-data and only indirectly perceive mind-independent objects. However, this contrast is unhelpful. As noted above, in 2.1, this seems like a purely linguistic issue, rather than a substantive theoretical one. If two theorists agree that the character of our experience is determined by sense-data, which are caused by mind-independent objects, then disagreement about whether or not we perceive sense-data seems like a debate about how to use the word “perceive.” Thus, we should be careful not to make too much of the claim that sense-data are the objects we perceive.²³ Moreover, understanding “direct” in this way renders most contemporary indirect realist views—indeed most, if not all, contemporary views of perception—as direct. Indirect realists like Sollberger argue that we perceive mind-independent

²³ In cases where it seems like this is not simply a terminological disagreement between different indirect realist sense-data theories, and where the claim that we perceive sense-data appears theoretically important, we may find that the sense-data theorist in question is a phenomenalist, rather than an indirect realist (e.g., Ayer, 1940).

objects. We may be only aware of such objects via being directly aware of mind-dependent ideas, but perception is a relation between us and mind-independent objects, not between us and ideas.

Disjunctive and non-disjunctive theories can therefore agree that we do not perceive mind-independent objects via perceiving something else. We can certainly say that these theories thus take us to perceive mind-independent objects directly, but the use of “directly” here seems needless. In contrast, labelling naïve realism a “direct” theory can help us make a useful distinction, as it highlights the fact that on such a view mind-independent objects are directly present to the mind in the sense that they are constituents of experiences and determine the character of perceptual experience. It is this second sense of directness that is fruitful to consider and pertains to the nature of perception.

We may think that the intentionalist will still be unhappy with their view being considered indirect. After all, intentionalists deny that what we are aware of in experience are mental objects or properties. Instead, they claim that we are aware of the properties of external objects. For instance, Harman argues that even if experience involves a mental picture, what one “is aware of is whatever is represented by that mental picture” (p. 36). He and Tye stress that the phenomenal properties that appear to us in experience—such as colours—are properties of mind-independent objects like trees and oceans. However, in doing so, they seem to employ “aware of” in a different way than others.

Indirect realists and critical realists take the phenomenal character of experience to be determined by mental properties; the phenomenal qualities that appear to us are properties of inner representations or states. It is in this sense that we are said to be aware of mental properties—it is mental properties that determine the character of our experience. And

intentionalists make the same sort of claim. According to intentionalism, the intentional properties of an experience determines its phenomenal character. We see this clearly in Harman and Tye, though Byrne is perhaps most explicit:

According to [intentionalists], the sensational component of a perceptual experience cannot vary independently of its intentional component: the phenomenal character of a perceptual experience is entirely determined by the experience's propositional content—that is, by what it represents. (2001, p. 199)

...there is a basic claim that all these philosophers [i.e. intentionalists] wish to defend. It is that the propositional content of perceptual experiences in a particular modality (for example, vision) *determines* their phenomenal character. (p. 204)

So the intentional properties determine the phenomenal character of experience. But intentional properties are properties of the mental state, not of the external objects. How the world appears to us—what appears to us in perception, the properties that appear to us—is determined by the representational mental state. In claiming we are aware of the properties of external objects, then, intentionalists seem to be expressing the idea that we perceive such objects. However, this need not place them at odds with other non-disjunctive theories. More importantly, since it is the intentional properties of an experience that determines its phenomenal character, intentionalism is an indirect theory in the sense I am concerned with here. It is by claiming that the intentional properties of an experience—and not mind-independent objects themselves—determine its

phenomenal character that the intentionalist can accept the common kind assumption (see, e.g., Fish, 2004; Crane, 2005).²⁴

If the intentionalist is going to say that we can be in the same intentional state whether we are veridically perceiving some object or having a hallucination, then their theory cannot be direct. Directness is a feature of theories like naïve realism, which take mind-independent objects to be constituents of our experiences, determining the character of our experiences. Thus, while the intentionalist can claim that their view is one on which we directly perceive external objects, it is nevertheless indirect in the relevant sense.

Thinking about non-veridical cases and the common kind assumption is not the only way to arrive at this distinction between direct and indirect views. If we instead start by thinking about the role of mental representation in perception we will find ourselves contemplating the same sense of indirectness.

Indirect realism, critical realism, and intentionalism take perception to involve mentally representing the world. On these views, when I perceive an object I am mentally representing it.

²⁴ At least part of the reason why intentionalists claim we are aware of the properties of mind-independent objects and deny we are aware of anything mind-dependent has been touched on in discussing the argument from transparency. If intentionalists understand being aware of something as being aware of it as a particular thing, then it is easy to see why they would deny we are aware of mental qualities. But, as noted above, in 3.2, indirect realists and critical realists do not argue that we are aware of mental qualities *as being mental qualities*. Harman seems to think of being aware of something in this way when he notes that when Eloise sees a tree, ...the colors she experiences are all experienced as features of the tree and its surroundings. None of them are experienced as intrinsic features of her experience. (p. 39)

However, the indirect realist and critical realist do not dispute this point. The issue is not how things appear to us, or whether colours appear to us as properties of this or that; but, rather, the issue is what determines the character of our experience, or in virtue of what things appear to us the way they do. It is therefore not surprising that intentionalists do not talk of our being aware of mind-dependent properties, even if only the intentional properties of our perceptual states. At one point Harman does note that Eloise “is aware only of the intentional or relational features of her experience” (p. 39). But, here, he seems to mean the point of view, or how the world is represented; this does not appear to be how he expresses the idea that the character of her experience is determined by its intentional properties.

What appears to me in experience depends on my representation of that object—that is, the character of my experience is determined by how I represent that object. This is what makes these views indirect. Mind-independent objects are not simply taken to appear to us in perception, and the character of experience is not taken to be determined by such objects and their properties. Rather, we represent such objects. Thus, we are not directly related to mind-independent objects in the way that we are on naïve realism.

It is certainly true that these three views have different ways of understanding the notion that perception involves mentally representing the world. Consider an indirect realist view like Sollberger's, on which mind-independent objects cause mind-dependent ideas. Such ideas possess phenomenal qualities, which we are aware of in experience, and represent the objects that cause them. Here, when I mentally represent the world an idea—a mental object—is generated. There is a mental representation that has phenomenal properties, and we can describe the representation as having certain properties.

For a critical realist like Coates, perceptual experiences are inner states possessing phenomenal properties, properties we are aware of. The inner—i.e. mental—states are caused by external objects, and their properties stand in a representational relationship to external objects.²⁵ On such a view, when I mentally represent the world there is no object generated which I am aware of. Instead, I am simply in a particular mental state that has phenomenal properties. Even though there is no mental object, we can describe the properties of this representational mental state.

For an intentionalist like Harman, perceptual experience is a matter of mentally representing the external world. Mentally representing the world involves causal interaction. As

²⁵ There is also the representation aspect, or representational content, associated with the conceptual component of the experience.

a result of such interaction, we end up in perceptual states. When I have an experience and mentally represent the world, I am in a mental state with certain representational properties. And it is the properties of this mental state that determine the character of my experience. It is stressed that on such a view there is no mental object we are aware of, but there is a mental state whose properties we can describe.

While the three approaches have a different way of articulating just how perception involves mentally representing the world, all take what appears to us in experience, or the phenomenal character of experience, to depend on the representational mental state a person is in and its properties. Such a mental state is taken to represent the world, but be distinct from the mind-independent objects it represents. Though all take mind-independent objects to play a causal role in the generation of such mental states, mind-independent objects are not constituents of the states. The views are therefore similarly indirect. Moreover, by examining how perception involves representation we arrive at the same sense of “indirect” discussed just above.

We have already seen that these three approaches accept the common kind assumption. Indeed, it is hard to see why anyone who thought perception involved mentally representing the world would reject the common kind assumption. Since representational mental states are distinct from the objects they represent, it is possible to be in the same state in the absence of the relevant object. Faced with illusions and hallucinations, and able to offer the same sort of account one gives for veridical cases, it seems natural to do so.

The question of whether or not perception involves mentally representing the world pits naïve realism against our three other approaches. The question cuts across the same lines as the common kind assumption and the direct/indirect distinction. It is, like the intimately related common kind question, a central question in debates about the nature of perception.

There are a few reasons, though, why we may want to focus our attention on the question of whether perception involves mentally representing the world.

First, doing so does not require that we consider the possibility of hallucinations that are indistinguishable from corresponding veridical experiences, nor hypothetical neurological interventions. Some naïve realists group illusions with hallucinations, but this approach is far from universal. Debates about the common kind assumption thus typically focus on hallucinations. This is what we find in the causal argument, where some hypothetical neurological intervention results in a subject experiencing a hallucination subjectively indistinguishable from a veridical experience that occurs when their brain is in the same state. Discussions of such hypothetical cases, however, risk shifting quickly into disputes over various principles relating to the nature of causation or obscure properties allegedly involved in the typing of mental states (see, e.g., Nudds, 2009).

Second, the key motivation for rejecting the common kind assumption seems to simply be the preservation of naïve realism, rather than any independent philosophical justification (Martin, 2004; Crane, 2005). If so, debate about the common kind assumption appears misplaced.

Third, focusing on mental representation gives us a way to understand the claim that two states are of the same fundamental kind. If perceptual experience involves mentally representing the world, this may explain why veridical and non-veridical cases should be understood as being of the same basic mental kind. Justifying the claim that two states are of the same kind requires deepening our understanding of the states in just the sort of way we must in order to articulate the way in which perception involves mental representation.

Fourth, whether or not perception involves mental representation is a question that allows us to better engage with empirical work on perception. Focusing on mental representation meshes with empirical work concerning perceptual processing and neural representations, as well as the philosophical literature discussing this work. Considerations related to typing mental states and whether we should explain different experiences in the same way will not be cast aside entirely. But thinking about things in terms of mental representation allows us to focus on individual cases of experience and how to explain them, as well as on how to understand perceptual processing as it actually takes place in the brain.

5. Naïve Realism versus Indirect Realism

In the previous section I argued that whether or not perception involves mental representation is a particularly important question to consider. In the next chapters I will look at whether recent empirical work on perception can help us answer this question. In what remains of this chapter I want to suggest that we can understand this question as constituting the disagreement between indirect realism and naïve realism.

Naïve realists deny that when I see my computer this involves my mentally representing a mind-independent object and that what my experience is like is determined by this representational state. If it therefore turns out that perception does involve mental representation in this way, then we must reject naïve realism.

For mental representation to be involved in this way is for there to be a mental state that is caused by and represents mind-independent objects, and for what appears to me in experience—the sensory qualities—to be determined by this mental state and its properties. This

is just what indirect realism accepts. Thus, Sollberger's indirect realism can be essentially understood as accepting that perceptual experience involves mental representation in this way.

It is true that Sollberger understands perceptual experience in terms of mind-dependent or mental objects, i.e. ideas. When I see my computer it is not simply that I am mentally representing the object, but there is a mental representation to speak of. My experience has certain sensory qualities because the mental representation—the idea—has certain properties. But there are two important things to note about the indirect realist characterizing experience in terms of mental objects. First, this seems like a natural way to think about experience given the role of representation. If I am mentally representing some external object, it seems natural to say that there is therefore a representation. Such a representation is certainly very different from external, physical objects, so I am not sure how much is added to our understanding of the representation simply by calling it a “mental object.” But there seems to be no harm in talking about a representation and its properties here.

Second, even though I think it is easiest to interpret things in terms of mental objects, this is not essential to the indirect realist view. As Sollberger notes:

Indirect realists have provided different ultimate analyses of sensory experiences that fall into two broad categories. Sensory experience is construed either in terms of sense data or else as the content of a state of sensing. According to the traditional and historically prior view, sensory of phenomenal qualities belong to entities called “sense data,” “sense,” “sensibilia” or ideas. For present purposes I will follow Locke's terminology and call them “Ideas.” (p. 815-6)

Although Sollberger concentrates on what he calls the “representative theory of perception,” which features ideas, we can view his key commitments as basically neutral between views on

which the character of experience is determined by the phenomenal qualities belonging to an inner object and where the character of experience is determined by the phenomenal qualities belonging to an inner state.

Of course, indirect realism is not alone in taking perceptual experience to involve mental representation. Critical realism also accepts that perceptual experience involves mental representation in the way described above. Likewise, intentionalism accepts that when I see my computer this involves my mentally representing a mind-independent object and that my experience is the way it is because of this mental state and its properties. However, according to critical realism perception also has an additional conceptual component. Understanding experience to comprise both a sensory and a conceptual component is significant for proponents of the view. Meanwhile, intentionalism adds that the intentional properties of an experience completely determine its phenomenal character. This distinctive commitment is an additional claim by the intentionalist, going beyond the more basic commitment to perception involving mental representation.²⁶

We can understand the three approaches, then, as being in the same family of views. Though critical realism and intentionalism are realist views that are indirect in the same way as indirect realism, the three are not identical. If it turns out that perception does involve mental representation—if our question is answered affirmatively—then we will have to look more closely at the issues that set the views apart. In short, adjudicating the disputes between these views takes us beyond our central question. But an affirmative answer to our question supports

²⁶ Note that this additional claim may make it difficult for the intentionalist to adequately explain the sensory qualities of experience. By claiming that the phenomenal character of an experience is exhausted by its intentional content, the intentionalist may leave us wondering how perception differs from thought—or, more precisely, how they can account for the difference between thought and perception given the commitments of their theory.

indirect realism. Whether it supports indirect realism in such a way that we should accept the theory as true depends on if we have other reasons not to do so, and on if the other views in the family are equally supported. Empirical support for one view will not necessarily be support for all of the views in the family—this depends on the arguments based on recent empirical work, and whether they are neutral between the views in the family or if they support some view(s) more than others.

Thus, the contrast between indirect realism and naïve realism is particularly revealing. The views stand on opposite sides of a crucial question, and progress can be made by asking whether recent empirical work at all bears on the disagreement between the two views.

Chapter 2

Recent Empirical Work and the Nature of Perception: A Review of the Literature

1. Introduction

In the next two chapters I consider the following question: Does recent empirical work give us reason to think that naïve realism is false or that indirect realism is correct? Recently, a few authors have argued that certain empirical work does one or both of these things (Smythies and Ramachandran, 1997; Smythies, 2005, 2009a, 2009b, 2011; Brown, 2008; Nanay, 2014; Pautz, 2016). I will argue that recent empirical work on perception does indeed give us reason to think that naïve realism is false and that indirect realism is correct.

In this chapter I review the literature that appeals to recent empirical work in order to argue against naïve realism. I will examine the different arguments that have been offered, as well as the empirical work that has featured in these arguments. I argue that the most promising argument is what I will call “the mismatch argument.”

In the subsequent chapter I examine a variety of empirical findings that can be used in the mismatch argument. This discussion will lead to the formulation of what I will call “the argument that mismatches are ubiquitous.” I will then evaluate these two arguments in detail and consider possible replies and objections, concluding that both arguments succeed in demonstrating that the empirical work in question suggests that indirect realism is correct.

2. Nanay on Multimodal Processing

Nanay argues that multimodal perception poses a problem for “anti-representationalist” theories of perception, like naïve realism. Specifically, he argues that “multimodal perception

seems to require matching two representations” (p. 46). Given that multimodal perception is pervasive, the naïve realist denial that perception involves mental representations appears to doom the theory.

As Nanay notes:

There is a lot of recent evidence that multimodal perception is the norm and not the exception—our sense modalities interact in a variety of ways. Information in one sense modality can influence the information processing in another sense modality at a very early stage of perceptual processing. (p. 45)

He cites ventriloquism as an example of such interaction. Here, vision influences audition and we hear voices as coming from a dummy rather than the ventriloquist. In presenting his argument, however, he focuses on a case where audition influences vision: the flash and beeps case. Here, an individual is presented with a flash of light accompanied by two beeps, and experiences two flashes (Shams et al., 2000).²⁷ A more well-known example of cross-modal influence is the McGurk effect (McGurk and MacDonald, 1976). In this study, when individuals viewed videos of people mouthing certain syllables that were dubbed with incongruent audio, they often reported hearing a different syllable altogether. For example, when lip movements for [ga] were accompanied with audio of [ba], individuals have the auditory experience of [da]. When eyes are closed, one clearly hears [ba]; conversely, muting the audio allows one to easily read the lips as mouthing [ga] (for discussion of such effects, see Bertelson and de Gelder, 2004).

Nanay is right to claim that multimodal perception is the norm. Countless experiments demonstrate the variety of ways in which different modalities may interact (for reviews, see

²⁷ The authors call this effect the “sound-induced flash illusion.” See Shams and Kim (2010) for discussion of this and other related effects.

Stein and Meredith, 1993; Driver and Spence, 2000; de Gelder and Bertelson, 2003; Bertelson and de Gelder, 2004; Calvert et al., 2004; Shams and Kim, 2010). Many of these involve dramatic examples of cross-modal influence—of the sort we see with ventriloquism, the flash and beeps case, and the McGurk effect—however, others demonstrate more subtle effects, such as how an unrelated auditory stimulus may enhance visual perception (Bolognini et al., 2005). These findings are good evidence that information from different modalities is integrated during perceptual processing; only if there is such integration can we see the sorts of effects uncovered in these experimental settings.²⁸ If the neural connections required for such interactions exist, then it is reasonable to suppose that this sort of integration occurs all of the time, even when no clear influence results.²⁹ Imaging studies and physiological investigations also reveal the connections between brain areas that process information from different modalities, as well as how certain neurons respond to stimuli from various modalities (see, for example, Stein and Meredith, 1993; de Gelder and Bertelson, 2003; Murray et al., 2016). Hence, there is now consensus that perception is multimodal and the integration of information from different modalities takes place (see, for example, O’Callaghan, 2012).

I will discuss multimodal perception in more detail in the next chapter. What matters for our present purposes is that multimodal integration takes place. In discussing Nanay’s argument I will, like him, simply focus on the flash and beeps case.

²⁸ We may notice similar effects in our daily lives—for example, how speech may be clearer if we can see the speaker’s mouth.

²⁹ Consider something like the McGurk study. It is hard to see how the visual and auditory information will be integrated *only* in the incongruent case, leading to the observed effect. The incongruence is only detectable once the information from the two modalities is integrated. The point here is simply that it is reasonable to suppose that such integration always occurs, even if we may only have behavioral evidence for it in certain cases. This is how research of psychological and neurological processes and functions typically operates; the experimental findings are taken to reveal the normal workings of our perceptual systems and processing.

In the flash and beeps case an individual is shown a single flash of light accompanied by two beeps. The auditory stimuli influences visual processing, resulting in the individual experiencing two flashes of light. In order for such influencing to occur, the information from vision and audition must be integrated during processing. Furthermore, the visual and auditory information carries information about the scene—for example, what was presented and where. As Nanay observes, “the multimodality of perception presupposes that information from two different sense modalities is unified in a shared framework” (p. 45-6). Without this, it would not be possible for information from different modalities to be integrated in a way that would result in the sorts of effects observed.³⁰

Nanay argues that it is easy for a theorist who understands perception to involve mental representation to explain the flash and beeps case. Vision represents the scene in a certain way, and audition represents it in a different way. The visual and auditory information are taken to come from the same source, so the visual system takes there to be two flashes—corresponding to the two beeps—rather than one.

However, he thinks that the theorist who denies that perception involves representation is in trouble.

To put it very simply, multimodal perception seems to require matching two representations, a visual and the auditory one. If we cannot talk about perceptual representation, how can we talk about what is being matched? The auditory sense modality gives us a soundscape and vision gives us a visual scene and our perceptual system puts the two

³⁰ In contrast, if the only sorts of effects observed were the strengthening or weakening of experience in a modality, it would be possible that the integration simply involved the sensory signal from one modality influencing the processing in another modality in virtue of its strength.

together. It is difficult to explain this without any appeal to representations. (p. 46)

Essentially, Nanay argues that we can only understand the processing involved in multimodal perception in terms of mental representations: multimodal perception involves the integration of information from different modalities, and we can only explain this by invoking mental representations. The problem for the naïve realist is that “the relation between the perceiver and the token perceived object that constitutes perception” on their theory, “seems to be the outcome of this process of unifying multimodal information,” which involves representations (p. 46).

Given that naïve realism denies that perception involves mental representations, Nanay’s claim is that the theory is unable to accommodate what amounts to normal perceptual phenomena.

The argument begins with recent empirical findings about multimodal interactions. Such findings are evidence that multimodal integration is the norm. The only way to explain this integration countenances mental representations, however this is inconsistent with naïve realism.

There are at least two ways for the naïve realist to respond to Nanay’s argument. First, the naïve realist could deny that mental representations must be invoked in order to explain multimodal integration. Second, the naïve realist could accept that representations are involved in the way Nanay claims, but argue that this is not inconsistent with naïve realism.

Nanay does not argue that in order to explain perceptual processing in general we must invoke mental representations. Rather, he argues that the processing involved in multimodal perception demands this. It is not entirely clear, however, why this sort of processing should pose a particular difficulty for the naïve realist. Nanay’s explanation of multimodal integration may seem perfectly compelling to someone who is disposed to accept that mental representations are involved in perception. But such a person would also be perfectly happy to view visual processing on its own in representational terms. If the naïve realist is able to, for example,

explain unimodal visual processing without any appeal to representations in the visual system, then it seems reasonable for them to ask why this more complex processing cannot be explained without representations. And since Nanay argues only that multimodal processing poses a problem, we may think he is willing to concede that the naïve realist can adequately explain unimodal processing.

However, we may think that the burden is not on Nanay to explain why this processing is distinctive, but is instead on the naïve realist to offer their own explanation of multimodal integration that does not feature mental representations. We could try to imagine what such an explanation would look like by looking at how naïve realists explain perceptual processing in general. The problem here is that naïve realists typically do not discuss perceptual processing. One presumes that the theory can accommodate the fact that perception involves the sort of processing revealed by scientific investigations throughout the twentieth century—if not, it is hard to understand how the theory could have any support. In the absence of such discussions, it is hard for me to imagine just how the naïve realist would explain multimodal integration.³¹ But it does not follow that no explanation may be forthcoming. Given what I have said here, I do not see why such an explanation could not be provided by the naïve realist.³²

A second way for the naïve realist to respond to Nanay's argument is to accept that representations are involved in the way he claims, but deny that this is inconsistent with naïve

³¹ Fish (2009, p. 134-40) briefly discusses how naïve realists may view perceptual processing at a very general level, but not at a level of detail required to help us understand how he may respond to the challenge here.

³² We may think the explanation would be that causal interaction between stimuli in the world and our sense organs initiates a number of steps of processing, and that at some point during this processing the information from different sources interacts, and then after some more processing we have perceptual experience. If given such an account, we may want to question whether it actually constitutes an explanation. However, we may want to question Nanay's account in just the same way. He claims that we should understand the information in various parts of our brains as representing the world in certain way, but is this really much more of an explanation?

realism. This is certainly a strategy that seems available to the naïve realist. They can claim that representations are involved at some stage of perceptual processing, but that this does not undermine their analysis of perceptual experience. After all, Nanay does not argue that we must invoke representations in order to explain the phenomenal character of perceptual experience. Therefore, the naïve realist may argue that Nanay's explanation of multimodal integration is perfectly consistent with their theory.

Nanay considers a reply along these lines. He considers the naïve realist arguing that there are sub-personal perceptual representations, but not any personal level representations. He takes this claim to be more or less equivalent to claiming that there are no representations involved in experience. He views this reply as inadequate because it renders naïve realism a theory of conscious perception; however, he argues elsewhere that not all perception is conscious (2014; Berger and Nanay, 2016).

The naïve realist may object to Nanay's response here on the grounds that it is question-begging. If the naïve realist either rejects the idea that there is unconscious perception or simply views their theory as an account of conscious perception, then they will be unmoved by Nanay's concerns. Whether or not one is sympathetic to the naïve realist here, I think it would be preferable for the debate about whether perception involves mental representations to be settled independently of the question of whether there is unconscious perception. For this reason, I do not find Nanay's response to the naïve realist manoeuvre compelling.

If we focus only the flash and beeps case, the naïve realist claim that their view of perceptual experience is unchallenged may strike us as puzzling. This is because of a feature of the case I have not focused much attention on: while the individual experiences two flashes, there is only one flash. In trying to think of how the naïve realist may explain the case without

invoking mental representations, we may have imagined them doing so in terms of objects or events in the world. However, they would face a problem here, given that individuals experience two flashes when there is only one.

Nanay, though, does not focus on this aspect of the flash and beeps case. And he does not argue that in order to explain the character of one's experience in such a case we must appeal to mental representations. I will come back to these sorts of multimodal interactions in the next section, as well as how we may point to such cases while making a different argument against naïve realism and in support of indirect realism. Here, I am focusing only on the argument Nanay does make. And his argument revolves around multimodal integration, which is not exclusive to non-veridical cases.³³

Nanay argues that multimodal integration gives us reason to think that naïve realism is false. However, his argument is far from decisive. In particular, it seems like naïve realists have two ways to avoid his conclusion. If they can offer an explanation of multimodal integration that does not invoke mental representations, they will show that Nanay is wrong in claiming that such integration can only be understood in terms of representations. As I say above, I do not see why we should consider such an explanation any more unlikely than an explanation of perceptual processing more generally. Second, they can simply maintain that we do not need to understand perceptual experience in terms of mental representations. Since Nanay does not dispute this

³³ Nanay's interest in the flash and beeps case is that it involves multimodal integration, not that it seems like the naïve realist may have trouble explaining the experience of two flashes given that there is only one flash. One reason why Nanay may not dwell on what individuals experience is because it would seem like the naïve realist, being a disjunctivist, will simply dismiss the case as non-veridical and so irrelevant. Because multimodal integration happens all of the time, the naïve realist cannot sidestep the issue in the same way. However, since it is not obvious how multimodal integration may relate to our experience in veridical cases, focusing on integration leaves open the second sort of reply discussed above.

point, there may not be any conflict between Nanay's explanation of multimodal integration and naïve realism. Given that both avenues appear viable, I will set aside Nanay's argument.

3. The Mismatch Argument

According to Smythies and Ramachandran, a recent binocular rivalry study demonstrates that naïve realism is false and that indirect realism is correct. What the empirical findings show, they argue, is that what appears to us in perceptual experience is determined by a representation generated by the brain.

In binocular rivalry studies different images are simultaneously presented to the two eyes. What individuals experience is not the two images superimposed on each other or mixed together in some way, but the images alternating, one after the other. Thus, if an image of a face is presented to one eye and an image of a house to the other eye, individuals will experience face and house percepts alternating every few seconds. The brain areas active during these different experiences can even be studied with functional imaging techniques, indicating which brain areas are involved in different sorts of perceptual experiences (Tong et al., 1998).

In the binocular rivalry study discussed by Smythies and Ramachandran, the images presented to the two eyes in the condition of interest were composites made by combining segments from two coherent images (Kovacs et al., 1996). Image A was a picture of chimpanzees and Image B was a picture of jungle foliage. When these are presented in binocular rivalry, individuals experience the two images alternating, as expected. Images C and D were patchworks, each made up of portions of Image A and Image B. The patchworks were constructed such that all of Image A and Image B were used, and the portions of those images appearing in the patchworks appeared in the spatial location they occupied in the original

images. When Image C and Image D are presented in binocular rivalry, individuals do not experience the two images alternating; instead, individuals experience Image A and Image B alternating.³⁴

Smythies and Ramachandran argue that this experiment provides empirical evidence that what they call the “Direct Realist” theory is false, and that it supports what they call the “Representative Theory.” I will first present the relevant passage from their article in full, then, after clarifying a few issues, construct the argument I think is contained in the passage.

After describing the condition of interest with Image C and Image D, they continue:

But what will theory predict if we expose one eye to picture C and the other to picture D? The Direct Realist theory would have to predict that the subject would see C and D in retinal rivalry: that is, after all, what is out there. However, that is not what happens. The subject sees A alternating with B. ...

The Representative Theory states that we do not see what actually is out there but what the brain computes is most probably out there. In these experiments, since the training of the brain’s computational networks have featured many unitary coherent pictures and very few incoherent patchworks of two coherent pictures, under the conditions of this experiment the brain restructures the two incoherent signals from the two retinae into the two coherent pictures that the subject actually sees. Thus the Direct Realist theory is refuted, for it would obviously be most implausible to suggest that we see only what the brain computes to be probably out there when looking at C and D, but not when we are

³⁴ The same thing happens when Image A and Image B are patterns of coloured dots or other stimuli, instead of pictures of things like monkeys and trees.

looking at A and B. Perception must depend on a brain mechanism that is unlikely to function in such an arbitrary fashion. (p. 438)

The first thing I want to note is their use of “see.” It is clear from the above that when Smythies and Ramachandran use “see” they mean visual phenomenology. That is, when they claim that “the subject sees A alternating with B,” they are making a claim about the subject’s experience and what it is like: they are claiming that what appears to the subject is a picture of chimpanzees alternating with a picture of foliage. They do not use “see” in the sense where the thing we see is whatever object we are perceptually related to. The indirect realist will claim that I see my computer, which is an external object, even if the phenomenal qualities I am aware of and that appear to me are mind-dependent. I have tried to use “see” and “perceive” only in this way. Thus, I would not say that the subject sees A alternating with B, but that they experience an image of chimpanzees alternating with an image of foliage, or that what appears to them in experience are these two images. That being said, I think that the way Smythies and Ramachandran use “see” is perfectly normal. If one were a subject in this sort of rivalry study, it would be perfectly natural to say: “I see the image of the monkeys alternating with the image of foliage.” Similarly, consider the flash and beeps case from the previous section. Here, it seems perfectly natural to say that individuals see two flashes, even though there is only one. Since there is only one flash, there is a sense in which one cannot see two flashes. However, we clearly also use “see” to simply mean our visual experience, whatever its relation to the external world. Used in this way, it is true that individuals see two flashes. So even though I think their use of “see” is innocuous, I will avoid using “see” in this way in order to avoid confusion. It is important to keep this all in mind, though, when trying to construct their argument.

Smythies and Ramachandran dub their target the “Direct Realist theory of perception.” This theory, they say, “states that, in vision, the visual field contains the physical object itself,

and thus the phenomenal object is identical to the physical object” (p. 437). Though some of the vocabulary may be idiosyncratic, it seems that they understand the Direct Realist view as one on which mind-independent objects are directly present to us and are constituents of our experiences. I take this to be captured by their claim that, on the theory, “the visual field contains the physical object itself.” This suggests to me the sort of directness we find on naïve realism, where physical objects are taken to be what appears to us. Hence, the “phenomenal object,” or what appears to us in experience, is just the physical object.³⁵ It would follow, then, that the phenomenal properties we experience are properties of physical objects.

It is reasonable, then, to think that the view they have in mind is naïve realism. At the very least, they have in mind an approach with certain core commitments, commitments which naïve realism has. Therefore, we can understand their argument as targeting naïve realism.

We may wonder whether naïve realism actually makes the prediction they attribute to the Direct Realist theory. It is certainly true that naïve realists claim that, in perception, we experience the mind-independent objects that are out there in the world. It would seem, then, that naïve realists would predict that when one eye is presented with Image C and the other with Image D, individuals will see both of them, or a mixture of them, or possibly both in retinal rivalry—but, certainly, not Image A and Image B alternating. However, naïve realists may not make this prediction because they may consider this case of rivalry an illusion—or, certainly, not a case of veridical perception. If so, naïve realism is not committed to individuals in this case experiencing the mind-independent objects that are out there. However, as we will see below,

³⁵ The expression “phenomenal object” may suggest they think experience must be construed in terms of objects, whether mind-independent external objects or mind-dependent objects like Söllberger’s ideas. However, I think it is fair to interpret Smythies and Ramachandran to simply mean the sensory qualities we are aware of in experience. This may be clearer from considering the Representative Theory, which I do next.

whether or not naïve realism makes a prediction about what our experience will be like in this specific rivalry case is irrelevant.

The Representative Theory, meanwhile:

states that the phenomenal object is a construct of the central nervous system and thus phenomenal objects are not identical to physical objects, but rather represent them, in which process a series of most complex neurocomputational mechanisms are involved that lie between the retinal image and the final construct—the phenomenal object in the visual field in consciousness. (p. 437)

Once again, their language may be idiosyncratic, but I think the core commitments of the Representative Theory are relatively clear, and so it is clear what sort of view they have in mind. The view adopts the sort of scientific causal picture discussed in the previous chapter, on which perception is a causal process that begins with physical objects interacting with our sense-organs and results in an inner state or object that represents but is distinct from the physical objects. Because Smythies and Ramachandran claim that the phenomenal object is a construct of the central nervous system—or, for simplicity, of the brain—we get a better idea of what they mean by “object” here. If the brain is constructing the phenomenal object, then we can understand it to be mental object. Since the phenomenal object represents physical objects, we can think of this as a mental representation.

But we don't necessarily have to think of the phenomenal object as a mental object. It seems that we can likewise understand the phenomenal object as what appears to us in experience, or the sensory qualities we are aware of in experience. What is essential is that these are generated, or constructed, by the brain. What distinguishes the view from the Direct Realist theory is that what we are directly aware of in experience are not physical objects, but something

mental, something generated by the brain—where what we are aware of, of course, is what determines the character of our experience. I think that either way of understanding “phenomenal object” is compatible with what Smythies and Ramachandran say. It may be easier to think of a mental representation generated by the brain as being a mental object, but I do not think it is necessary to do so.

What is clear is that on the Representative Theory “the phenomenal object in the visual field in consciousness” is what appears to us in experience, and that the sensory qualities that appear to us are determined by an inner state or mental object that represents the world. The physical object is not directly present to the mind, and since the phenomenal object is not identical to the physical object, the physical object and its properties are not constituents of our experience. The Representative Theory, then, is a theory on which perception involves mentally representing the world and what our experience is like is determined by the representational state (whether thought of as a brain state or a mental state). The theory is simply indirect realism.

We are now in a position to articulate the argument offered by Smythies and Ramachandran. In essence, what they point to is how in the binocular rivalry study in question our experience differs dramatically from what is external to us. There is a mismatch between what one’s perceptual experience is like and the actual physical objects before us in the world. The question is then how to explain this experience, and whether that explanation is consistent with naïve realism or indirect realism. Their argument, which I will call the “mismatch argument,” is the following:

(i) In the rivalry study, an individual is presented with two patchwork images. The mind-independent objects causally interacting with the individual’s sense-organs are the jumbled images C and D.

(ii) The individual experiences two coherent images alternating. Their perceptual experience is of chimpanzees alternating with foliage and they do not experience jumbled images.

(iii) We can explain what the individual experiences if we suppose that the brain constructs a representation that determines the character of their experience. The individual is aware of this representation. We cannot explain the character of the individual's experience simply by pointing to the external objects, since their experience, or what appears to them, differs in the way it does from the external objects. In this case, the individual is not immediately aware of the external objects.

(iv) Perceptual experience depends on brain mechanisms that always function in the same way. If the character of our experience in certain instances is determined by a representation constructed by the brain, then the character of our experience is always determined by a representation constructed by the brain.

(v) Therefore, in perceptual experience we are aware of a mental representation generated by the brain; we are never aware of mind-independent objects and their properties.

From this conclusion it follows that naïve realism—which says that perception involves the direct awareness of mind-independent objects, which determine the character of our experience—is false. Rather, perception is a matter of mentally representing the world, and the character of our experience is determined by the representation. Indirect realism is vindicated.

I have used the same terminology here as in the previous chapter. What we are aware of in perception is what determines the character of our experience. Smythies and Ramachandran do not use “aware of,” but they think that how things appear to us in experience is due to the representation constructed by the brain and not the external objects. In (v), we may want to say that we are never *directly* aware of mind-independent objects—doing so changes nothing.

The first two premises of the argument are straightforward. They simply describe the relevant features of the situation and highlight the mismatch between the objects external to the individual and the individual's experience.

In (iii), the issue is how to best explain the individual's experience in the rivalry situation. The argument does not hinge on what, if any, predictions different theories make about experience in this situation. It is about explaining the experience, and whether an explanation consistent with a given theory is adequate or even available. We can understand their claim about what theory predicts as concerning explanation. Naïve realists claim that the character of experience is determined by external objects. In this case, experience is of coherent images, but the external objects are patchworks. We cannot explain what is experienced, then, in terms of the external objects and their properties.

It is also important to note that the idea that the brain constructs a representation does not materialize out of nowhere, simply in order to explain this finding. This idea, and the theory of perception supported by Smythies and Ramachandran, finds favour against a background of scientific work and theorizing. The claim that the brain constructs a representation that determines the character of our experience is made within the context of a scientific understanding of perception and the brain, an understanding which is based in a wide range of experimental work.³⁶

This scientific understanding of perception and the brain is also behind (iv). The idea here is that perceptual processing is fundamentally the same whether or not we are in a binocular

³⁶ It is not this one study alone that leads Smythies and Ramachandran to reject naïve realism and support indirect realism. They present this study as a refutation of naïve realism, but they do so as supporters of indirect realism addressing proponents of a naïve realist approach they see as misguided. As they note, there "have been many experiments in the past on which the current theoretical attitude of visual scientists is based" (p. 437).

rivalry study. The way information is processed in the brain does not vary depending on whether or not there is a mismatch between our experience and what is external to us. The crucial aspect of that processing, for the argument, is that the brain constructs, or generates, a representation. Since perception depends on brain mechanisms that always function in the same way, the brain always generates such a representation.

As Smythies and Ramachandran note, it would be implausible to suggest that the brain constructs a representation when we look at images C and D, but not when we look at A and B. The brain has access only to its sensory inputs, and does not know whether we are presented with C and D or A and B. What happens in the rivalry study is that the brain interprets the inputs as resulting from A and B and takes the noise—i.e. the jumbled nature of the inputs—to be due to the perceptual system rather than the external objects. The brain does not generate a representation because it somehow decides that what is external to us is unacceptable. The representation is generated in the course of normal perceptual processing—it differs from what is external to us in the rivalry study because the brain supposes that we are actually presented with A and B.

This point relates to the discussion of probability and the brain computing “what is most probably out there.” According to Smythies and Ramachandran, when the brain constructs a representation it generates a representation of what it takes to be external to us given the sensory inputs. Throughout our lives we encounter many more coherent scenes than patchworks. The sensory inputs that result from the coherent scenes arrive in the context of further inputs, including following our navigating our environment, which itself is guided by the representations constructed by our perceptual systems. These inputs provide feedback regarding the adequacy of the representations constructed on the basis of prior inputs. This is the “training of the brain’s

computational networks” they mention. As a result, when we are presented with C and D, our brains determine that we are presented with A and B and that the input is simply very noisy.³⁷

The same mechanisms are at work, though, when we are presented with A and B. Here, the brain determines that we are presented with A and B (rather than, for example, being presented with C and D but there being lots of noise).

The mismatch argument bears something of a resemblance to the argument from illusion presented in the previous chapter.³⁸ Here is that argument:

- (i) In an illusory experience it seems to one that something has a property, *F*, which the mind-independent object perceived lacks. Though the wall seems red to me, it is actually white, not red.
- (ii) When it seems to one that something has a property, *F*, then there is something of which one is aware that has that property. Traditionally, sense-data are invoked here, as the objects we are aware of. Thus, what I am aware of is a red sense-datum.
- (iii) Since the mind-independent object in question does not have the property, *F*, then in illusory experiences one is not aware of the mind-independent object.
- (iv) The same account of experience must apply to both veridical and illusory experiences.
- (v) Therefore, one is never perceptually aware of mind-independent objects. What one is aware of in experience are sense-data.

³⁷ The brain is simply wrong in this case. Also, it is worth noting that all sensory input is noisy to a certain degree. I will return to this issue in the next chapter. I will also discuss the relation between what the brain constructs and probability in my final chapter.

³⁸ Given the similarity between the arguments from illusion and hallucination, the same resemblance holds regarding the latter. I will focus on the argument from illusion here for two reasons. First, for ease of discussion. Second, because the rivalry study seems to be an instance of illusion and not hallucination—on this point, see discussion of mismatches in the next chapter.

Both arguments begin with a situation where an individual's experience differs from the mind-independent object(s) presented to them. For example, in the relevant experience what appears to the individual has certain properties that the object lacks. In order to explain the experience, the individual is said to be aware of something besides the mind-independent object(s). It is then claimed that the same account should be given for all perceptual experiences.

However, there are some important differences between the argument from illusion and the mismatch argument. First, the mismatch argument is not cast in terms of illusory and veridical experiences. While the focus is on particular perceptual experiences, the argument does not presuppose a clear theoretical distinction between cases of illusion and veridical perception (and even hallucination). Nor does the argument depend on certain experiences being classified as illusions or hallucinations.³⁹ Second, premise (iv) of the mismatch argument is supported by a deep scientific understanding of perceptual processing and the brain mechanisms involved. It is not simply claimed, as in the argument from illusion, that the same account should be given for veridical and illusory experience, either because these are indistinguishable or because they form a continuous stream. Rather, the premise in the mismatch argument that plays the same sort of role as the spreading premise in the argument from illusion—seeking to generalize a claim about some experiences to all experiences—is grounded in our understanding of the actual mechanisms that are involved in perceptual processing. The reason why it is claimed that the same sort of account will apply to all experiences is because in all cases the same well-understood neural mechanisms are involved in the same way. Third, the mismatch argument does not invoke sense-data, nor is it claimed that in experience we must be aware of some object. Its premise (iii)

³⁹ As we will see, some discussions of mismatches do involve reference to “illusion,” “misperception,” or “illusory” aspects of an experience. However, even in such cases there is some question about just how we ought to characterize illusions and the like.

concerns how we are to explain the relevant experience, and doing so within the context of our understanding of the brain.

Faced with the mismatch argument, naïve realists seem to have two options: they can dispute (iii) or deny (iv). The first strategy would require that the naïve realist explain the character of our perceptual experience in the rivalry study without invoking a mental representation, but simply by appealing to the external objects. We have already seen how the naïve realist seems to be in trouble here. Examining whether the naïve realist can successfully do this will involve looking at how naïve realists have tried to account for illusions. The second strategy would be to deny that perceptual experience depends on brain mechanisms that always function in the same way. This would be to appeal to disjunctivism. Whether or not this manoeuvre is viable, and how the naïve realist could motivate it, becomes a crucial question.

The naïve realist may be tempted to simply dismiss the binocular rivalry study as a highly abnormal situation that we should not pay too much attention to. However, there are a number of findings that can be used to run the mismatch argument.

Elsewhere, Smythies points to other empirical findings to support the conclusion that the brain constructs our experience. Studies revealing spatial and temporal filling-in offer other instances of a mismatch between what we experience and the mind-independent world.

Spatial filling-in is “the interpolation of missing information across visual space” (Weil and Rees, 2011, p. 41). Information may be missing due to the lack of photoreceptors where the optic nerve exits the eye or to damage to the retina or early visual areas resulting in a scotoma. While scotomata have a variety of causes, we all have such a blind spot where the optic nerve exits the retina. However, we do not experience blank areas in our visual field. Instead, filling-in

occurs and our visual experience is unbroken. Because there is never any visual information coming from such blind spots, filling-in is a ubiquitous process.⁴⁰

While filling-in of the blind spots and scotomata happens instantly and regardless of the type of visual stimuli, there are various other kinds of filling-in, as well as a variety of proposed mechanisms (for reviews see Komatsu, 2006; Weil and Rees, 2011). I will discuss spatial filling-in in more detail in the following chapter. For our present purposes, it is sufficient to note that spatial filling-in typically involves filling-in information from the surround—i.e. the area around the blind spot—such that we may experience some property in a part of the visual field despite its not being present in the world (Komatsu, 2006). That is, we will have cases where there is a mismatch between our experience and what is external to us at the blind spot. A simple example would be the following: if I close one eye and look at a beige wall with a red spot, and the spot happens to fall in my blind spot, my experience will simply be of a spot-less beige wall.

Temporal filling-in is related to saccadic suppression, the phenomenon whereby during saccades information from the eyes is suppressed. Despite the lack of visual input, there is no temporal gap in our visual perceptual experience. This is because the brain extends our visual experience backwards in time, to cover the period lacking visual information (Yarrow et al., 2001). Saccades are so brief that any discrepancy is likely to go unnoticed, but as a result of temporal filling-in there may be a mismatch between our experience and what is out there in the world at the corresponding time. One instance where we do notice something is when we saccade to a clock with a second hand, and for a brief moment the clock appears static, the second hand sticking for a moment longer than it subsequently does. This occurs because “the

⁴⁰ In normal binocular vision information from one eye may compensate for the lack of information from the other. But this does not matter for our present purposes. It is easy to demonstrate with monocular viewing how filling-in may result in our failing to see something.

brain extends the percept of the saccadic target backwards in time to just before the onset of the saccade...[to] ‘fill in’ the perceptual ‘gap’ during saccadic suppression” (p. 302).

We can thus present the mismatch argument by starting with a mismatch due to filling-in, rather than the binocular rivalry case. Consider a case where there is a mismatch between what one’s perceptual experience is like and the actual physical objects before us in the world, such that the object has some property or feature that we do not experience (or vice-versa). The first two premises of the argument would therefore be:

(i) In such a case, an individual is presented with an object having certain properties, including something on its surface at a certain location. The mind-independent object causally interacting with the individual’s sense-organs is a certain way.

(ii) What the individual experiences does not have the relevant property in the relevant location, but instead has some other property in its place. Their perceptual experience is of an object where the properties of the area surrounding the blind-spot are filled-in.

Premises (iii), (iv), and (v) would then follow, as above.

Instances of multimodal interaction like the flash and beeps case and the McGurk effect can also be used in the mismatch argument. In the previous section, I noted an aspect of the flash and beeps case that Nanay does not dwell on: while the individual experiences two flashes, there is only one flash. There is a clear mismatch between what the individual in the study experiences and the stimulus in the world. Similarly, with the McGurk effect there is a mismatch between our auditory experience of [da] and the audio being played, which is the sound [ba]. In order to explain these experiences and their phenomenal properties, it seems like we must countenance mental representations.

The mismatch argument with the flash and beeps case would be:

- (i) In the flash and beeps study, an individual is presented with a single flash of light.
- (ii) The individual experiences two flashes of light.

Premises (iii), (iv), and (v) would once again follow, as above.⁴¹

I will come back to the mismatch argument in the next chapter and assess it in detail. The key questions will concern premises (iii) and (iv), so whether we should explain certain perceptual experiences with mental representations, and whether it is correct to think that the brain always generates such representations. I have already noted some of the reasoning supporting these premises, as well as how naïve realists may try to respond. Before examining these issues more closely, I will look at two other arguments from the recent literature.

4. Brown's Case for Indirect Realism

Brown also argues that recent empirical work supports indirect realism. In presenting what he calls “a new argument from illusions” he points to recent findings regarding illusory contours and colour perception, as well as to the detailed knowledge we have of visual processing from scientific investigations. His argument is similar to the mismatch argument.

Though Brown considers his argument an argument from illusion and introduces the empirical findings that motivate his argument as illusions, it is not clear that all of the findings he alludes to should be considered illusions. This is true even if we understand “illusion” in the way Brown does, as “any qualitative mismatch” between our perceptual experience and the external objects that cause our experience, on which no deception is required (p. 46). I will examine each example he discusses in turn.

⁴¹ Thus, the mechanisms and processing that always function in the same way, alluded to in (iv), do not concern multimodal integration but the generation of a representation by the brain. It is true that integration always occurs, but that is not the relevant processing here.

The first “class of illusions” Brown points to are illusory contours. Illusory contours are the experience of lines, forms, or volumes in the absence of such features in the stimuli (Murray and Herrmann, 2013). For example, when viewing a Kanizsa figure made up of three pacman shapes we experience a triangle framed by the three shapes. We have the experience of a triangle here even though there is no change in luminance along the areas of the stimuli corresponding to the triangle’s sides. Many different stimuli may induce illusory contours, and a variety of mechanisms have been proposed to explain the related phenomena. These experimental findings shed light on the normal functioning of the visual system and the processes underlying our perceptual experience. What is clear is that the visual systems plays an active role in creating illusory contours.

Given that illusory contours involve a mismatch between our experience and the objects or stimuli external to us, Brown is right to consider such phenomena illusions.

Next, Brown discusses colour. Recent empirical findings show how our experience of colour depends on various factors in addition to the light reflected by a portion of some surface, including light reflected by the surround and overall illumination (Purves and Lotto, 2003; Allred and Olkkonen, 2013). Memory and previously viewed stimuli can also impact colour perception (Olkkonen and Allred, 2014; Olkkonen et al., 2014). Because of the myriad factors that contribute to our colour experience, the same stimuli can be accompanied by different colour experiences in different contexts. Studies of colour perception help reveal how the brain processes ambiguous sensory input in an effort to achieve constancy, however imperfectly (Olkkonen et al., 2016).

Does this empirical work reveal mismatches between our perceptual experiences and the external objects that cause our experiences? The point being made here is not that our experience

has certain properties that physical objects entirely lack.⁴² Rather, the point is that we can have two different experiences of a target stimuli without changing the stimuli, and we can have the same colour experience with different target stimuli placed in different contexts. The same sort of independence of our experience from what is external to us that we see with the other mismatches seems to be present here. What is evident, and what Brown highlights, is the brain's role is disentangling the sensory input, and how our experience is shaped by this processing, rather than being straightforwardly shaped by the external stimuli and incoming light. Colour is a complicated issue, and I will return to it in the context of mismatches in the next chapter. But, I think the empirical work supports the idea that there are at least some qualitative mismatches between our experience and the external objects that cause our experience when it comes to colour. At the very least, colour constancy is imperfect.⁴³

Finally, Brown notes how felt heat and cold depends on the relative temperature of our bodies. Regarding this, he says:

...we recognize that felt heat and cold are not on a par either ontologically or causally. An item feels warm to my hand when there is

⁴² The allegation is not that all colour experience involves properties that physical objects lack—this is not an argument from secondary qualities.

⁴³ Brown also points to the case of experiencing black when a portion of the stimuli does not reflect any light. About this, he says:

Phenomenologically black is on par with the other colors. When [an experience] includes black areas, nothing in my experience indicates that my perception of black is brought about in a different way from my perception of other colors. ...Yet the external conditions that produce black are different from those for other colors. (p. 47)

It is not clear to me why Brown thinks experiencing black involves a mismatch, i.e. is an illusion in the sense he defines. He notes that experiencing black is similar to experiencing other colors even though they may differ causally. The point, I take it, is not simply that the causes are different, but that the experience of black may be caused in a manner that differs from the manner in which all other colour experiences are caused. But it is not clear why we should think that the different causal origin should be evident in the experience, as he seems to suggest. Nor is it clear to me that this is a case of illusion in Brown's sense. The experience of black when no light is reflected is not a mismatch in the sense I have outlined above.

heat flow from that item to my hand; it feels cold when heat flows in the reverse direction. (p. 48)

It is unclear why exactly Brown thinks our experiences of heat and cold are therefore illusory. Whether this is the case would seem to largely depend on what the sensations of heat and cold are sensations of, an issue which he does not explore.⁴⁴ In any case, his contention that such experiences involves mismatches is unsupported. Therefore, I do not think this third “finding” should be considered a case of illusion.

Still, this leaves Brown with mismatches relating to illusory contours and colour for the argument he articulates. Before examining his argument, a couple of bits of terminology require clarification. An “external arrangement” is the collection of things—whether objects or events—that cause a perceptual experience, where “external” means only that these things exist independently of the relevant sense organ. And we can understand the “perceptual display” to include whatever we are immediately aware of in perceptual experience.

Brown argues for indirect realism in two steps. The first step is:

...in normal perception the perceptual display we are aware of is caused by an external arrangement and (let us grant) there is a one-one mapping between areas of the display and areas on the surface of its external cause. But there are often substantial qualitative differences between properties of these corresponding areas, and qualitative

⁴⁴ He goes on to note that the “special causal conditions” of an experience of cold “are not apparent to introspection,” but he does not explain what its being apparent or not has to do with whether there is a qualitative mismatch between an experience and the objects causing it. Furthermore, this issue appears secondary to the question of what we should view such sensations as sensations of. His mention of what is apparent to introspection is similar to the point he makes about experiencing black (see fn. 43). I find the point puzzling in both cases, and it leads me to wonder if he may understand “mismatch” in a much broader sense than I do. However, I do not think this issue is too significant, given that the first two things he points to are mismatches, as I understand mismatches.

difference implies numerical difference. To this extent the perceptual display is not wholly numerically identical with the external arrangement that causes it. One might maintain that the perceptual display is a heterogeneous compound of elements of the external arrangement and elements generated by our perceptual system. But given the coherent integration of the various elements of the perceptual display, we get a more intelligible picture if we consider the entire perceptual display to be a brain-construct that is numerically distinct from the external arrangement involved in its causation. (p. 48-9)

He notes that this is an argument to the best explanation that “provides *initial grounds* for holding that perceptual displays are numerically distinct from the external arrangements that cause them” (p. 49).

The second step is to look at how perceptual displays are generated, i.e. to look at our scientific understanding of perceptual processing. Here, Brown points to some of what we know about the early stages of visual processing, focusing on the retina. In essence, he presents the sort of scientific causal picture discussed in the previous chapter, then claims:

There is no reason for holding that [the outcome of this processing] is, somehow, numerically identical with one contributor to the causal process that produced it. This conclusion would hold even if there were qualitative identity between an external arrangement and a perceptual display. (p. 53)

Thus, even though Brown thinks the scientific causal picture suggests indirect realism, he points to recent empirical findings revealing mismatches in order to support the conclusion that the perceptual display is distinct from the external arrangement, i.e. that external objects are not

immediately present to us in experience. He thinks that the combination of these two steps constitutes an argument for indirect realism. I will now examine his argument more closely.

I take “normal perception” simply to mean normal cases of perceptual experience, at the exclusion of perhaps only hallucinations. The first part of his argument is the following:

- (i) In normal perception the perceptual display is caused by an external arrangement, i.e. external objects and events.
- (ii) Suppose a one-to-one mapping between areas of the display and areas in the external arrangement.
- (iii) Often, there are substantial qualitative differences between properties of these corresponding areas. That is, there are often mismatches between our perceptual experience and the external objects that cause our experience.
- (iv) Qualitative difference implies numerical difference.
- (v) Therefore, the perceptual display is not identical to the external arrangement that causes it.
- (vi) Given the coherent integration of the perceptual display, it makes most sense to think of it as completely distinct from the external arrangement that causes it.

The first two premises seem uncontroversial, and would happily be accepted by his opponent. The recent empirical work Brown discusses supports (iii). These findings support the claim that there are sometimes, or even often, mismatches between our experience and the external objects that cause our experience. In such cases, there is qualitative difference between the perceptual display and the external arrangement. And in such cases it follows, from (iv), that the perceptual display is not identical to the external arrangement. We can understand this conclusion, so (v), as opposing the naïve realist line that what is present to us in experience are mind-independent objects, and that the character of our experience is determined by such

objects. At the very least, Brown has argued that some parts of our experience involves the awareness of something mind-dependent, that some features are due not to external objects but to our perceptual system. Then, given that our perceptual experience is coherently integrated, it makes sense to think of it in its entirety as involving the awareness of something other than the mind-independent objects that may cause it. He further supports this last claim by pointing to our scientific understanding of perceptual processing.

However, note that (iv), (v), and (vi) only concern the cases where there is a mismatch, i.e. the cases of illusion he points to in order to make his argument from illusion. He claims that the first part of his argument gives us initial grounds for thinking that perceptual displays are numerically distinct from the external arrangements that cause them. But the key step en route to this conclusion is qualitative difference, which he claims implies numerical difference. Of course, we only have qualitative difference in cases where there is a mismatch. This point would seem to be hardly worth making; in fact, given that Brown takes the time to point to different findings regarding mismatches, it is unsurprising that these play an important role in his argument. But, notice that in summarizing the first part of his argument Brown talks about perceptual displays and external arrangements in general, without distinguishing between those instances where there is a mismatch from those where there is not.

In articulating the first step in his argument, Brown focuses on the relationship between a perceptual display and the external arrangement that causes it. He then notes that there are often mismatches, and that in such a case the display is not wholly identical to the external arrangement. The rest of the argument focuses on that perceptual display in question, one qualitatively different from its external arrangement; it plainly focuses on a single experience, or a single perceptual display.

In response to the first step of Brown's argument, the naïve realist may be happy to concede his conclusion insofar as it concerns cases of illusion. However, the naïve realist, being a disjunctivist, could maintain that this does not in any way impact their view of perception. The naïve realist analysis, after all, is taken to apply to veridical cases, and Brown's conclusion that the perceptual display is distinct from the external arrangement would seem to only apply only to illusions. The naïve realist can therefore still claim that in cases where there is no qualitative difference we have no reason to think that the perceptual display is distinct from the external arrangement.

Since Brown presents his argument as an argument from illusion, one would think that he includes a spreading premise, or some consideration that is supposed to get us from what is the case in the instance of illusions to what is the case in all instances of perceptual experience. In the traditional argument from illusion, it is typically claimed that the same account should be given for illusory and veridical cases because these are indistinguishable or form a continuous stream. It may therefore be tempting to read his rejection of the idea that the perceptual display is a "heterogeneous compound" and the related (vi) as being such a claim. However, he is clearly here talking about a single perceptual display, rather than the relationship between perceptual displays.⁴⁵

What about the second step: is this the spreading premise? As we have seen, the second step is simply pointing to the scientific causal picture of perception. Brown thinks that if we look at how perceptual displays are actually generated, we will be inclined to accept his claim, made

⁴⁵ Though we may want to read "the perceptual display" in the second half of the lengthy quotation above as denoting perceptual experience in general, and so his point about heterogeneity to simply be the continuous stream claim, I think this would be a mistake. Not only does he seem to clearly be talking about a single perceptual experience, but elsewhere he talks about "perceptual displays"—plural—which would not make sense if "the entire perceptual display" simply meant the entirety of our experience.

on the basis of the first part of his argument, that perceptual displays are numerically distinct from the external arrangements that cause them. The problem here is that he has only given us reason to accept this conclusion when it comes to cases where there are mismatches. The “initial grounds” for thinking that perceptual displays are distinct from external arrangements are qualitative differences, which are only present in some cases. It would seem, then, that the second step simply supplements his argument about those cases and is not a spreading premise.⁴⁶

Regardless of how Brown presents the second step, we may think that the scientific understanding of perceptual processing he points to supports the claim that perceptual processing always functions in the same way. I think it is right that the scientific picture suggests this. The processing he describes is the same in all cases, whether or not there is a mismatch. Thus, if as a result of perceptual processing we take “the entire perceptual display to be a brain-construct” in some cases, we should take the perceptual display to be such in all cases. Admittedly, Brown does not discuss the latter stages of perceptual processing that would include the construction of a representation we are aware of—that is, the brain-construct he argues explains the illusory cases. But, given that he argues there is something constructed in some cases, and that the scientific understanding of perceptual processing he points to supports the idea that perceptual

⁴⁶ We may wonder whether this is the right way to understand the second step. After all, Brown claims that the second step alone suggests that perceptual displays are numerically distinct from the external arrangements that cause them. He claims this is so even in the absence of qualitative difference (p. 52-3). I find this remark puzzling in light of the first part of his argument, which relies on such qualitative differences. Given that in the first part it is only this qualitative difference that gives us reason to think that the perceptual display is distinct from the arrangement, then it is not clear why we should think the scientific causal picture suggests the same conclusion in the absence of any qualitative difference. Moreover, if the scientific causal picture does suggest this, it would seem he does not need the first part of his argument—or, rather than considering them two parts of a single argument they should be seen as two distinct arguments. Note that even if we think of the second step as somewhat independent of the first, it still is not a spreading premise, as it does not include the claim that the same account that is given for illusory cases—the sort of account he offers in the first step—should be given for veridical cases.

processing always functions in the same way, the second step would seem to support the conclusion that there is always something constructed by the brain. Therefore, even if he does not explicitly present the second step of his argument as helping him extend his conclusions about illusory cases to all perceptual experiences, I think he could appeal to it in order to do that work.

The question of whether or not Brown's argument includes something analogous to the spreading premise is directly related to the question of how similar his argument is to the mismatch argument. Though Brown points to different empirical findings from Smythies and Ramachandran, his argument is more or less the first three premises of the mismatch argument. Both arguments point to a discrepancy in some cases between our experience and what is external to us, and are concerned with offering the best explanation for our experience in light of that. Like Smythies and Ramachandran, Brown argues that what we are aware of in experience is constructed by the brain. And Brown's claim that the perceptual display is not identical to the external arrangement that causes it echoes Smythies and Ramachandran claiming that the phenomenal object is not identical to the physical object.

However, the mismatch argument goes beyond Brown, at least explicitly, with:

(iv) Perceptual experience depends on brain mechanisms that always function in the same way. If the character of our experience in certain instances is determined by a representation constructed by the brain, then the character of our experience is always determined by a representation constructed by the brain.

I certainly think that Brown would accept this premise. He may even take it to be implied by what he does say. So either we can view Brown as offering an argument roughly equivalent to the mismatch argument, or we can view him as offering a proper part of that argument. In either

case, nothing will be lost if we focus our attention only on the mismatch argument and how the naïve realist may try to respond to it. Moving forward, I will therefore focus on the mismatch argument—which is rather clear and straightforward—and leave Brown’s argument aside.

5. Pautz on Bad External Correlation

In the preceding sections we have seen different attempts to argue, based on recent empirical work, that naïve realism is false and indirect realism correct. In thinking about whether these arguments are successful, we may find ourselves considering the following question: How can empirical findings bear on the dispute between these two theories—that is, what are the different possible things empirical work may show us that would be relevant to the debate? Any answer to this question will be only preliminary, since we might not be able to imagine how certain empirical work can inform a debate until a certain sort of empirical work appears on the scene. But if we think about how empirical findings may give us reason to reject naïve realism or how empirical findings may support indirect realism, two main ways come to mind.

First, empirical investigations of perceptual processing may reveal the complexity of perceptual processing or the different kinds of processing that take place in our perceptual systems, where these are incompatible with naïve realism. Nanay pursues this strategy, alleging that the processing involved in multimodal perception poses a problem for naïve realism. As we saw, however, it is not clear why the complexity of the processing leading to our perceptual experience would require that we understand such processing in a way that is incompatible with naïve realism. For the facts about perceptual processing to be incompatible with naïve realism would seem to require that naïve realists make specific claims about perceptual processing, and not simply about the nature of perceptual experience. Nanay argues that the naïve realist’s

rejection of mental representations extends beyond perceptual experience to include perceptual processing in general. Whether or not he is right here, his argument is unsuccessful because he has not given us reason to think that the complexity of the processing involved in multimodal perception requires that we understand such processing in terms of mental representations.

Nanay's argument, though, points to a particular way we may think empirical findings about processing could bear on the debate: if empirical work shows that this processing includes the construction of mental representations. But why would empirical findings give us reason to think that mental representations are constructed in the course of our perceptual processing? This brings us to the second way empirical findings may come into play.

The naïve realist wants to explain our perceptual experience and its nature in terms of mind-independent objects and events, they want to explain the character of our experience in terms of the external world. The indirect realist wants to do so in terms of the mind and the brain. That our perceptual experience is closely related to our brain activity is not something that is disputed by any party in this debate—both sides agree that our brain activity is essential to our perceptual experience and that different experiences correlate with different patterns of activity. There is less agreement about how our experience is related to the external world. The indirect realist thinks our experiences are caused by and represent the external world. But the naïve realist thinks the relationship between our experience and the world is much tighter, and that external objects and events are directly present to us in experience, that they are constituents of our experience, and that they determine the character of our experience. Empirical findings may be relevant to the dispute if they somehow drive a wedge between experience and the external world by showing that the relationship is not so tight, is not as the naïve realist claims. If empirical findings give us reason to think that we cannot explain the character of our experiences

as the naïve realist would like to, or if empirical findings suggest that in order to explain what our experiences are like we must invoke mental representations, this would pose a problem for naïve realism. These two things go hand-in-hand, since one reason we may think a representation is involved is in order to explain an experience that cannot be explained by the external objects present; meanwhile, if some experience cannot be explained by some external objects present, pointing to a representation generated by the brain will offer a way to explain the experience.

We have already seen one way that empirical findings seem to drive a wedge between experience and the external world. Mismatches are cases where what we experience differs from what is out there in the world, external to us. According to the mismatch argument, this difference means that we must invoke mental representations in order to explain our experiences. Although he argues in a different manner than what we see in the mismatch argument, Pautz also tries to drive a wedge between experience and the world by pointing to recent empirical findings. I will claim, however, that his argument is unsuccessful.

Pautz argues that findings from neuroimaging studies, together with considerations about the properties of perceptual stimuli, give us reason to reject naïve realism. The problem with naïve realism, he contends, is that “it violates *internal dependence*: the empirically determined role of the *internal processing* of the brain in shaping phenomenal character” (p. 23). We can see this, moreover, by focusing only on what he calls “normal” cases of perception, so not illusions or hallucinations.

Naïve realists claim that the character of our experience is determined by external objects and their properties, which we are immediately aware of in experience. Though they concede that the brain has a role in perception, they think that external objects have objective sensible properties and that “when the brain responds to these objective sensible properties in the

biologically normal way, this *enables* the mind to ‘reach out’ and become acquainted with them” (p. 24). However, if we look at similarity relations between experiences, neural activity, and the physical properties of objects, Pautz thinks we will see that:

Looking for the basis of qualitative character in the external world is looking for it in the wrong place; we have to be looking rather at the brain. (p. 25)

First, Pautz points to colour experience. Consider my looking at a blue ball, purple grapes, and a green leaf. Pautz claims that “the blue appearance of the ball resembles the purple appearance of the grapes much more than the green appearance of the leaf” (p. 25). However, if we look at the reflectance properties of the three things we will find that reflectance of the ball does not resemble the reflectance properties of the grapes more than the reflectance properties of the leaf; if anything, the ball resembles the leaf more so in this respect.⁴⁷ So whereas my experience of the ball is more similar, in terms of colour, to my experience of the grapes, the reflectance properties of the ball are more similar to the reflectance properties of the leaf. Pautz calls this “bad external correlation”—there is a bad correlation between the similarity relations among my experiences and the similarity relations among the relevant properties of the external objects.

In contrast, if we look at brain activity, as revealed by neuroimaging, we find the same similarity relations as we do with the experiences. Imaging of the brain areas (e.g. V4) involved in colour perception demonstrates that similar colour experiences involve similar patterns of activity (Brouwer and Heeger, 2013). Thus, the pattern of brain activity correlated with the

⁴⁷ In his paper, the reflectance properties of the objects are depicted with graphs displaying the ratio of light reflected at various wavelengths by the three things (figures are from MacAdam, 1985). The shape of the curve for the blue ball resembles that of the green leaf more so than it does that of the purple grapes.

experience of the blue ball will resemble the pattern of brain activity correlated with the experience of the purple grapes more so than it will resemble the pattern of brain activity correlated with the experience of the green leaf. This gives us what Pautz calls “good internal correlation.”

Second, he discusses our odour experiences of three different chemicals. Two of the chemicals, citral and R-limonene, smell like citrus. The third, R-carvone, has a minty smell. However, the chemical structure of R-limonene resembles the chemical structure of R-carvone and not that of citral. Neuroimaging studies, meanwhile, reveal that the brain activity correlated with our experience of R-limonene resembles the brain activity correlated with our experience of citral, more so than it does the brain activity correlated with our experience of R-carvone. In general, imaging studies show that brain activity in odour areas of the brain will be similar for stimuli rated as having similar odour quality (Howard et al., 2009; Margot, 2009). Once again, therefore, we have bad external correlation and good internal correlation.

Third, Pautz considers auditory experience. Voice onset time is the amount of time between the opening of the mouth and vocal fold vibration. Pautz points to a study involving stimuli that will elicit an experience of [da] or [ta], depending on voice onset time. But whereas voice onset time varies continuously, the shift between the experience of [da] and [ta] is categorical—we hear [da] with a voice onset time of up to 30 ms, but after 40 ms we hear [ta]. Meanwhile, neuroimaging work shows that there is a categorical change in related brain activity corresponding to our auditory experience (Chang et al., 2010).

Because our experience and the corresponding neural activity display a categorical shift, but the voice onset time varies continuously, we once again have good internal correlation and bad external correlation when it comes to similarity relations. We can see this if we compare

three points in time. With a voice onset time of 30 ms our experience is [da], at 45 ms our experience is [ta], and at 80 ms our experience is [ta]. Our experience at 45 ms is more similar to our experience at 80 ms than it is to our experience at 30 ms—as is our brain activity in the relevant brain areas. However, a voice onset time of 45 ms is nearer to, or more similar to, a voice onset time of 30 ms than it is to one of 80 ms.⁴⁸

The findings from different modalities indicate that the similarity relations between our experiences are mirrored by the similarity relations between our corresponding neural activity, but not by the similarity relations between the objective properties of physical objects and events that cause those experiences.⁴⁹ We can order our experiences according to how much they resemble each other—Pautz calls this their “resemblance-order.” We can do the same for the brain activity correlated with these experiences, as well as the properties of physical objects that cause these experiences. Pautz’s contention is that naïve realists cannot explain the resemblance-order of our experiences, because this differs from the resemblance-order of the relevant objective properties. Naïve realists may claim that “sensible properties are *irreducible* objective properties of things that are *grounded in*...physical properties,” but “this requires *irregular grounding*: totally unsystematic and arbitrary grounding connections” (p. 28). Since we should prefer theories with systematic grounding relations, we should reject naïve realism.

⁴⁸ While it may seem like we can say there is bad external correlation simply because the voice onset time can vary continuously but our experience is either/or, demonstrating the lack of correlation requires comparing three points. If we look only at 30 ms and 80 ms, we have unanimous responses of [da] and [ta], respectively. By considering the profile of responses at a third point in time the curve between 30 ms and 80 ms begins to take on a shape that differs from what we would have with a continuous shift between the two.

⁴⁹ Though Pautz only discusses a few experiences and the corresponding brain activity and physical properties, it is reasonable to think that further investigation would reveal more relations of the sort he uncovers.

It is important to note that Pautz does not argue that naïve realists are unable to explain any particular experience by appealing to the properties of external objects. Given what he has argued, the naïve realist can point to the reflectance properties of the ball to explain the qualitative character of my experience of blue. Likewise, the naïve realist can point to the reflectance properties of the leaf to explain the qualitative character of my experience of green. And the naïve realist can point to the chemical structure of R-carvone to explain my minty odour experience. The problem for the naïve realist is that whereas the chemical structure of R-carvone is similar to the chemical structure of R-limonene, my odour experience caused by R-limonene is similar to that caused by citral, but not at all like that caused by R-carvone. It is the resemblance-orders (i.e. similarity relations) among experiences that naïve realists are, according to Pautz, unable to adequately explain. Meanwhile, the explanation for the resemblance-orders among our experiences seems to be brain activity.

So that is Pautz's argument. Now, does he succeed in driving a wedge between our experience and the external world, and showing that in order to explain the qualitative character of our experience we must look at the brain rather than the external world?

One thing a naïve realist may do in response to the argument is question the resemblance-orders that Pautz determines. There are three different sorts of resemblance-orders to consider, and the naïve realist may have concerns about each.

Take the relative similarity of different sensible qualities. Is it clear that the blue appearance of a ball is more similar to the purple appearance of grapes than it is to the green appearance of a leaf? Is it not possible for people to disagree about which colours are more or less similar, or to think that blue is more similar to purple in some respects, yet more similar to green in other respects? Meanwhile, we may think that even if people's judgments about whether

a stimulus is [da] or [ta] shifts categorically at a certain voice onset time, the stimuli around this time may be less clear than at the extreme ends of the continuum. People may readily identify a stimulus as [da] even if they find certain stimuli in this group are more or less clear than others, or that they do not all sound identical even if they all sound like [da]. Though I think the naïve realist would be right to question these resemblance-orderings, I am not sure how much traction they will get here.

Their case seems to be stronger regarding the properties of physical objects. There are two different questions they may ask here. First, what is the relevant property to consider? Second, are there different ways to represent the relevant property, where these different ways of doing so will have different resemblance-orders? Take the odour example. Here, Pautz notes how the chemical structure of R-limonene resembles that of R-carvone more so than it does that of citral. Looking at diagrams of their structures, one must agree. But even if we think that odour is related to the chemical properties of objects, we may wonder whether the relevant physical property that causes our experiences is the entire chemical structure of the molecule—as we may depict with a molecular model—rather something else, or something more specific. And even if we think we should consider the chemical structure, it may not be the case that sketching a molecular model is the only, much less the best, way to do so.

Regarding colour, we may have no doubt that reflectance is the right property to consider. However, there are different ways to represent, or characterize, reflectance. Pautz presents three curves depicting the reflectance properties of three objects (from MacAdam, 1985). In terms of the general shape of the curves along the x-axis, from 400 to 700 nm, it is true that the shape of the curve for the blue ball resembles the shape of the curve for the green leaf more so than it does the shape of the curve for the purple grapes: the curves for the ball and leaf slope upwards to a

peak, then curve downwards until rising just slightly before 700 nm; meanwhile, the curve for the grapes rises and falls like the other two, but then begins to rise steeply around 550 nm and continues to rise until 700 nm. However, the entire curve for the leaf is much lower on the y-axis than that of the ball—where the value on the y-axis is “reflectance” of the surface, or the “ratio of the amount reflected divided by the amount incident” (MacAdam, 1985, p. 3). In fact, the curve for the ball is in the same range along the y-axis as much of the curve for the grapes. And, between 450 and 500 nm, where both crest before falling, the curves for the ball and grapes are at roughly the same point on the y-axis, while the curve for the leaf is much lower. We may very well think, then, that the reflectance of the ball actually resembles that of grapes more so than it resembles that of the leaf if our main concern is around the 450 and 500 nm range and the amount (or ratio) of light of these wavelength that is reflected. The point is that there are different ways to compare and contrast the curves, while we also may have empirically-informed reasons to focus on particular portions of the curves more so than others if we are interested in characterizing the reflectance properties of the objects.

The naïve realist may have the best case when it comes to the neuroimaging data and the alleged similarity relations here. The issue here is not whether we think that similar experiences will be correlated with similar patterns of brain activity, but how exactly we judge patterns of brain activity to be similar. Here, Pautz’s use of neuroimaging data seems dubious. We can certainly rely on imaging data to reveal which brain areas are more active in different conditions. But the resemblance-orders Pautz is concerned with depends on our comparing patterns of activity within certain small neural regions in the way we would visually compare ink patterns on a page, then determining in which cases brain activity is more similar simply by considering the spatial pattern of activations within these small neural regions, across subjects. Not only may

there be other, more appropriate, ways to assess the relative similarity of patterns of neural activity, but there are also the questions of fineness of grain and inter-subject differences. It seems reasonable to worry about concluding too much on the basis of the sorts of apparent similarities in patterns of activity Pautz points to.

In summary, it seems like the naïve realist has grounds for questioning the various resemblance-orders that are at the heart of Pautz's argument.

But even if the naïve realist is willing to bite the bullet and accept that sensible properties and the relevant physical properties fall into different resemblance-orders, we may not think this is a problem for their view. Accepting that these resemblance-orders differ means accepting the sort of irregular grounding Pautz alleges. Brewer accepts the charge of irregular grounding and points to other apparent cases of irregular grounding, arguing that if we accept those cases we should not reject naïve realism.⁵⁰ In response, Pautz claims that even if we are to accept some cases of “arbitrary, unsystematic grounding,” “we surely should *prefer* theories with more systematic grounding connections,” and so should prefer an alternative to naïve realism that explains the character of experience in terms of the brain (p. 29). The point certainly sounds compelling, and the charge of arbitrary and unsystematic grounding connections sounds damning. However, if we consider just what the naïve realist is allegedly unable to explain, and where exactly the grounding is irregular, I think we are right to wonder whether these are good reasons to reject the view.

Pautz does not argue that, for example, the naïve realist cannot explain the character of my experience of blue by pointing to the reflectance properties of the ball. When he says there is poor external correlation, this does not mean that particular experiences are poorly correlated

⁵⁰ Pautz mentions Brewer's response and notes that it was offered in person.

with the relevant physical properties of external objects. He does not show that some physical property can lead to different experiences on different occasions, nor that the same experience can be caused by different physical properties. His argument is not that the relationship between physical properties and our experiences are arbitrary or unsystematic, nor that if we think of the qualitative character of an experience as being grounded in the properties of external objects, that any of these connections are irregular or haphazard. The charge is not that the naïve realist is unable to explain the phenomenal character of any experience by pointing to the properties of external objects. Rather, the alleged problem is that the naïve realist cannot explain certain similarity relations without tolerating irregular grounding. But the naïve realist may simply reply that they are under no obligation to explain why sensible properties and the relevant physical properties fall into different resemblance-orders, and that we should not be concerned with irregular grounding here. And, I think we should be sympathetic to this reply. After all, the naïve realist is trying to explain our perceptual experience and its nature in terms of the external world. Once they do so, there is no need for them to say more; there is nothing more that they need to explain. The naïve realist can maintain that the connections between particular physical properties and corresponding experiences are systematic and regular—Pautz does not cast doubt on the strength of these connections.

Furthermore, I see no reason to think that the irregular grounding Pautz points to calls their explanation into question. For it to do so would require that the different resemblance-orders undermined the plausibility of the connection between the external world and our experience. But it is not clear to me why we would expect all resemblance-orders between various sorts of things to be uniform, such that we completely avoid irregular grounding of the sort Pautz describes. While I certainly expect a tight relationship between the brain and our

experience, such that patterns of brain activity are correlated with experiences, I am not sure whether we should expect that there will also be good correlations between different similarity relations we could tease out. And even if we do want to understand why the resemblance-orders in question differ, we need not think it is the job of a theory of the nature of perception to help us here.

Thus, I do not think Pautz succeeds in driving a wedge between experience and the external world, such that the naïve realist account of perception is called into question. His argument, and the empirical work he points to, does not give us a compelling reason to doubt that the naïve realist can explain the phenomenal character of our experiences in terms of mind-independent objects and their properties.

Chapter 3

Two Arguments Against Naïve Realism

1. Introduction

A review of the literature uncovers one promising argument that uses recent empirical work to argue that naïve realism is false and indirect realism is correct: the mismatch argument, which we find in Smythies and Ramachandran and, possibly, in Brown. The argument is:

- (i) In the rivalry study, an individual is presented with two patchwork images. The mind-independent objects causally interacting with the individual's sense-organs are the jumbled images C and D.
- (ii) The individual experiences two coherent images alternating. Their perceptual experience is of chimpanzees alternating with foliage and they do not experience jumbled images.
- (iii) We can explain what the individual experiences if we suppose that the brain constructs a representation that determines the character of their experience. The individual is aware of this representation. We cannot explain the character of the individual's experience simply by pointing to the external objects, since their experience, or what appears to them, differs in the way it does from the external objects. In this case, the individual is not immediately aware of the external objects.
- (iv) Perceptual experience depends on brain mechanisms that always function in the same way. If the character of our experience in certain instances is determined by a representation constructed by the brain, then the character of our experience is always determined by a representation constructed by the brain.

(v) Therefore, in perceptual experience we are aware of a mental representation generated by the brain; we are never aware of mind-independent objects and their properties.

The argument gets off the ground with a mismatch, an instance where what a person experiences differs from what is external to them. Any mismatch will do here. In addition to the binocular rivalry study with patchwork images, we have already seen how mismatches can result from spatial and temporal filling-in, multimodal interactions, illusory contours, and colour perception. The argument concerns the question of how we should explain these experiences, and the proposal in (iii) is that the explanation invoke mental representations. From there, the reasoning is that if there is a representation generated by the brain in these cases, then there always is. These two steps are the key points in the argument and it is regarding (iii) and (iv) that we can expect pushback. In order to determine if the argument is successful, we will need to see if we are justified in explaining these experiences with mental representations and if we are correct in thinking that perceptual processing always operates in the same way.

In this chapter, I will closely evaluate the mismatch argument and consider various responses to it. I will argue that none of the responses are effective. Before doing so, however, I will take a closer look at the empirical work relating to mismatches. This discussion will make plain that far from being exceptional occurrences elicited in the confines of the lab, mismatches are the norm in our day-to-day lives. The ubiquity of mismatches will allow me to offer a second argument in support of indirect realism: the argument that mismatches are ubiquitous. Via these two arguments, I will conclude that recent empirical work suggests that perception involves mentally representing the world, such that the phenomenal character of our perceptual experience is determined by a representation. That is, this work supports indirect realism, and compels us to reject naïve realism.

2. Mismatches

Mismatches are interesting at least in part because they seem to pose a problem for anyone who wants to say that we directly experience physical objects or are directly aware of external objects and their properties. In these cases objects appear to us to have certain properties or features they actually lack or we have the experience of things that are not actually out there in the world. Such a discrepancy calls out for an explanation, an account of what is going on such that external objects and events lead to these experiences.

That these experiences make such a demand points to a second reason mismatches are interesting: they are a window into the functioning of our perceptual systems.⁵¹ By studying situations where mismatches occur we can uncover fundamental principles of perceptual processing. Hence, there is a lot of scientific research relating to mismatches. And it is often findings from such work that are behind the common refrains that “perception is active” and “perception is construction.”

It may be tempting to simply think of mismatches as cases of illusion, but I will stay away from calling mismatches either “illusions” or “hallucinations.” “Illusion” is a notoriously slippery expression, and some even question whether there is a clear category of phenomena to which it can be applied (Calabi, 2012; Rogers, 2014). It is sometimes used, as by Brown, to simply mean a qualitative mismatch, but illusions are sometimes thought to involve deception or to require a salient or dramatic effect. Moreover, they may be understood in purely perceptual terms or as involving a cognitive or epistemic component.

⁵¹ We would want to understand the mind and perception even in the absence of mismatches, and there is a lot we can learn about perception without focusing on such cases. However, these phenomena seem especially to pique our interest and in many cases offer a unique opportunity to understand certain perceptual processes.

Meanwhile, “hallucination” is no less troublesome. Philosophers often consider hallucinations to be cases of perceptual experience where there is no perception and are preoccupied with (possibly only hypothetical) total hallucinations, episodes which are indistinguishable from cases of veridical perception but where no objects are present. Of course, even in the paradigmatic drug-fuelled or psychosis-induced hallucinations there will be things perceived in conjunction with things hallucinated, and the latter may even be intimately related to the former. It also seems perfectly acceptable to talk of elements or aspects of an experience that are hallucinatory, and for these to be divorced from delusion.

More mundane cases reveal how the line between illusion and hallucination may be unclear. Consider the flash and beeps case from the previous chapter. Here, we experience something that is not there, i.e. a second flash. However, the authors call it the “sound-induced flash illusion.”⁵² The McGurk effect and certain cases of filling-in also seem to blur the boundary between illusion and hallucination.

I think that some mismatches will naturally seem to be illusions, others will seem more like hallucinations, and in many cases neither label will seem appropriate (of course, this may depend on one’s preferred definitions). I will stick with “mismatch,” and for now not concern myself with the question of whether a particular mismatch—much less mismatches in general—is an illusion, hallucination, or something else. If I talk about an “illusory” or “hallucinatory” element of some experience, I will be using these in a rough, pre-theoretic sense.

⁵² Whether this case strikes us more as an illusion or hallucination may depend on the timing of the flashes, so whether we experience two flashes during a period of time in which the single flash extends, or if the two flashes extend beyond the time over which there is the actual flash. These things, however, can be hard to measure precisely. We also may have questions about properties and objects—e.g. is it a property of the visual stimuli that it is one flash as opposed to two? Finally, it is not clear whether it makes sense to think of the beeps being experienced as flashes, or if this would stretch the boundaries of misperceiving.

In this section I will take a closer look at the empirical work relating to mismatches. In the previous chapter I briefly discussed multimodal perception, binocular rivalry, spatial and temporal filling-in, illusory contours, and colour perception. Here, I examine the relevant empirical work in more detail and consider other findings of the same kind. Examining this work will help us gain a better understanding of perceptual processing across different modalities and shed light on how common mismatches—often slight, unnoticed, and of little or no consequence—are in our daily lives.

2.1. Binocular Rivalry

Binocular rivalry studies can uncover how perception functions, even if the mismatches we find here will not occur outside of the laboratory setting. Typical rivalry studies involve presenting images of different sorts of things simultaneously to the two eyes. For example, a face may be presented to one eye and a house to the other. Functional imaging can reveal which brain areas are more active prior to and during these different conscious experiences. This work can help shed light on which brain areas are involved in processing different kinds of visual information, as well as which areas are involved in conscious awareness.

That we experience the two images alternating rather than as superimposed or side-by-side also tells us that visual experience is not simply the culmination of a process whereby input to the two eyes is combined in a manner that reproduces the stimuli presented to us. Even though the situation is highly artificial, our experiences here point to the sorts of processing occurring between retinal stimulation and conscious experience.

The rivalry study involving patchwork images can tell us even more about this processing. That our experience alternates between stable images in typical rivalry studies shows

a bias towards coherence. That we experience stable images alternating even when we are presented with patchworks is stronger evidence of such a bias. Furthermore, the study with patchwork images begins to point to the sorts of principles that may be guiding the various stages of processing that take us from input to experience. The images presented to the two eyes—whether patchworks of monkeys and foliage, or of patterns—are not blended haphazardly, but in a specific way. We can ask why the images are combined in that precise way. The hypothesis that this is related to prior sensory inputs makes sense. Hence, we start to see talk of the brain “interpreting” what is external to us, or “assembling” things, or computing what is “most probably” out there (Tong et al., 1998; Kovacs et al., 1996; Smythies and Ramachandran, 1997). The process will be shaped by prior sensory inputs, including inputs following actions guided by the impression of the world created on the basis of immediately prior inputs—i.e. feedback.

2.2. Multimodal Perception

There are a variety of different things one may mean by “multimodal perception”: one may be pointing to the fact that our conscious experience includes qualities from different modalities appearing together; one may be referring to the fact that information from different modalities is integrated during processing and that the different modalities are not causally or functionally encapsulated; one may be concerned with cases where we perceive the same object via two different modalities simultaneously, like when I see my hand on my computer and feel the keys beneath my fingers; or, one may even be referring to the awareness of certain features which, it is thought, we can only have via the coordination of multiple modalities (O’Callaghan, 2017). Typically, at least in the scientific literature, when people say that perception is multimodal they mean that information in different modalities is not processed independently of

information in the others, but that there is integration between modalities. This is what I will mean, too, when claiming that perception is multimodal. Perception is multimodal in the sense that perceptual processing involves the sort of multimodal integration briefly discussed in the previous chapter. It is this integration that makes interactions between the modalities possible, where information from one modality influences processing in another.

That such integration occurs is supported by a wealth of empirical work detailing the variety of ways different modalities may interact, as well as by work revealing the various connections between brain areas that process information from different modalities and the activity of multisensory neurons (see, for example, Stein and Meredith, 1993; Driver and Spence 2000; de Gelder and Bertelson, 2003; Murray et al., 2016). What this research demonstrates is that information from different modalities is brought together in such a way that allows for coordination across modalities. It is not simply that the processing in one modality can influence the processing in another, but that the information in one can influence the information in the other. For example, vision and audition can both map space, but for the spatial information in the visual system to be coordinated with spatial information in the auditory system requires some sort of shared spatial framework (for discussion, see Driver and Spence, 1998, 2004).

Experimental investigations have uncovered a variety of interactions between modalities. Ventriloquism reveals how vision can influence the perception of sound, moving where we experience a sound coming from closer to the visual stimuli (Bertelson and de Gelder, 2004). In the flash and beeps case, audition influences vision, and we experience two flashes when two beeps are presented with a single flash (Shams et al., 2000). The effect has been shown to also occur in reverse, when two flashes presented with a single beep lead to the experience of a single flash (Watkins et al., 2007). The interaction between vision and audition need not involve the

clear dominance of one modality. For example, in certain instances of the McGurk effect—e.g. when lip movements for [ga] are accompanied with audio of [ba]—the resulting auditory experience is of a different syllable altogether. Though interactions between vision and audition have been the most extensively studied, tactile stimulation has also been shown to influence vision (Shams and Kim, 2010), and audition has been shown to influence touch (Jousmäki and Hari, 1998) as well as experiences related to taste (Spence, 2015).⁵³

But not all crossmodal interactions lead to mismatches. Such interactions can result in the enhancement of a signal—as in when an auditory stimulus leads to the detection of an otherwise below threshold visual stimulus (Bolognini et al., 2005)—or improved accuracy—as in when vision and audition interact to improve spatial perception (Shams and Kim, 2010). This is not terribly surprising if multimodal perception is adaptive. In fact, coordination between modalities can contribute to the accuracy of our experience in a variety of ways. Since it has better resolution in the spatial domain, vision can improve the spatial perception of auditory stimuli.⁵⁴ In contrast, audition excels in the temporal domain, thus enabling auditory information to improve the visual perception of events close together in time (Hairston et al., 2006). Meanwhile, vision and audition work together to improve speech perception (Navarra et al., 2012) and to

⁵³ For discussion of additional effects and interactions, see Stein and Meredith (1993), Driver and Spence (2000), de Gelder and Bertelson (2003), Bertelson and de Gelder (2004), Shams and Kim (2010), and O’Callaghan (2012).

⁵⁴ It was previously thought that vision dominated crossmodal interactions, however it is now clear that interactions vary according to a number of factors. Whether a given modality may dominate in a particular instance can depend on the quality of the inputs, their timing, and even the modality’s appropriateness to the task at hand (Bertelson and de Gelder, 2004). For example, while vision often dominates spatial and motion perception, visual experience can be influenced by other modalities even here, which is not surprising given that motion can be difficult to perceive visually in certain contexts—e.g. darkness—in which case “sound and touch can convey information that can contribute to a more accurate perception of motion” (Shams and Kim, 2010, p. 272; also, see discussion in O’Callaghan, 2016).

achieve subjective simultaneity of physical events with both visual and auditory aspects (Fujisaki et al., 2004).

The notion that crossmodal interactions could render perception more accurate than if the modalities were causally and functionally encapsulated may be puzzling if our understanding of perceptual processing is overly simplistic or based solely on introspecting on our conscious experience. We may imagine that there are external objects and events that interact with our sense organs, which transmit the relevant information to our brains, thus allowing us to experience these spatiotemporal objects and events. With such a picture in mind it may be unclear how, for example, auditory information could make my visual perception more accurate—I already have the visual input and all the visual information I need, so how could auditory information be relevant?

Shams and Kim (2010) help shed light on this matter while discussing illusions resulting from crossmodal influences on vision:

Why should the visual system be allowed to be misled by other modalities? This kind of crossmodal interaction appears to be non-adaptive. Do these interactions represent a suboptimality in the human nervous system or are there advantages that justify having such interactions?

Intuitively, it can be seen that if there are two sensory measurements (e.g., auditory and visual) available about an environmental variable (e.g., the timing of an event), then given that sensory measurements are always noisy, it would be beneficial to combine the two measurements to obtain a more informed estimate of the environmental variable. More formally, it can be shown that if there are two noisy observations of the

same variable, if both observations are unbiased estimators, then integrating the two measurements can result in a more precise estimate.
(p. 279)

The key, here, is that not only do different modalities rely on different types of signals that may carry different sorts of information from the world, but there is noise. As de Gelder and Bertelson (2003) note, there is:

...moment-to-moment variability, or noise, in either impinging stimulation or subsequent processing. This noise usually affects one modality specifically and thus produces transient crossmodal incongruence. Combining information across modalities can bring compensation for these disturbances. (p. 460)

Aside from noise due to the incoming signals themselves (in part do to our own movements) and neural processing, there are temporal discrepancies that arise due to the disparate location of sense organs receiving information from the same source and different timescales for different neural processes (see, for example, Fujisaki et al., 2004). For instance, if two objects collide to my left and make a sound, we have to consider the time it takes for light to reach my retinas, for sound to reach each of my ears, and for the different sorts of brain activations to occur. In order for things to be appropriately bound in experience, the brain must take the different sensory signals to be caused by the same object or event.

The picture that emerges from this work on multimodal perception is somewhat complex. It is important to at least consider external objects and events, inputs to and information in perceptual processing areas in the brain, and our perceptual experiences. Because the inputs to the brain and the information at different stages of processing can be noisy, information from other modalities helps improve the accuracy of our experiences. This happens via interactions

between the modalities, when information in one modality influences that in another. Sometimes this influencing will operate mainly in one direction, but it is often bidirectional, helping to coordinate or calibrate neural signals. As a result, crossmodal interactions will often lead to our experiences more closely resembling the objects and events that cause them. Thus, multimodal perception is adaptive.⁵⁵

Where does that leave us regarding mismatches due to multimodal interactions? We have seen that there is multimodal integration and that interaction between the modalities is pervasive. We have also reviewed some of the findings revealing the sorts mismatches that may result. But these mismatches occur in the lab, where stimuli are manipulated to achieve effects. What about in daily life—how common are such mismatches here?

It is reasonable to think that mismatches due to multimodal interactions occur frequently in our day-to-day lives. While it is true that the laboratory setting involves the artificial manipulation of stimuli, the setting is also artificially simple and contains significantly less stimuli than one encounters even in the more controlled environments in the world, not to mention walking through an urban or natural landscape.⁵⁶ When we consider how we are bombarded with sensory stimuli as we get on in the world and how noisy these signals are it is easy to understand how mismatches can result from crossmodal interactions. Of course, these

⁵⁵ That multimodal perception is adaptive does not mean it results in experiences that are accurate in all instances or in all ways. However, one would think that, as compared to other available alternatives, it results in more or less accurate experiences, more often than not, at time-scales and spatial fineness of grain that matter for us.

⁵⁶ Moreover, the manipulation of stimuli in the lab may even allow this controlled setting to approach the real world in certain respects.

One major instrument of [multisensory integration] research has been the experimental conflict situation, in which two modalities receive incongruent data regarding one particular aspect of the environment. Contrary to a frequent misconception, such conflicts occur not only in the laboratory, but also in the natural environment, and artificial experimental conflicts can be seen as simulations of natural ones. (de Gelder and Bertelson, 2003, p. 460)

mismatches need not be as drastic as what we see in the flash and beeps case or with the McGurk effect. They may simply be a matter of something being slightly displaced in space or time, or there may be a relatively minor mismatch with an aspect of some visual, auditory, tactile, olfactory, or gustatory experience. Given the diversity of effects observed and the ease with which they can be elicited, we can expect mismatches due to multimodal interactions to be common even though multimodal perception is adaptive and renders our perception more accurate than it would be in the absence of such integration and interaction. Though the brain may rightly take multiple incoming signals to be related or to come from the same source, it can also mistakenly take unrelated signals to be related in this way or fail to associate related stimuli.

That such mismatches go unnoticed is unsurprising. Unless it is pointed out to us or conflicts with our subsequent experiences in an obvious way, we will not notice there is a mismatch as we go about our lives. This is particularly true if the mismatches are minor, insignificant instances within the flow of our perceptual experience. But even drastic illusions—for example, effects similar to what we get with the flash and beeps case—will likely go unnoticed. Another feature of the lab and the controlled setting is that discrepancies are closely measured, mismatches uncovered and trials repeated, and subjects debriefed. But when a naïve subject experiences the McGurk effect nothing appears amiss. That the experience involves a mismatch is not something we notice from introspecting on our perceptual experience. However, the abundant empirical findings concerning multimodal perception reveal the extent of such phenomena. And this work helps us see that given the flow of experience and how our perception of the world changes as we move amongst the objects in it, if some aspect of an experience involves a mismatch, this is likely to go unnoticed and be inconsequential.

2.3. Saccadic Suppression and Temporal Filling-In

From the first person perspective it seems like we open our eyes and the world simply presents itself to us—we open the shutters and there it is. But the reality is much more complicated. Our visual experience is constructed on the basis of information from our retinas, but while our visual experience when our eyes are open is stable and uninterrupted, retinal inputs are anything but. As a result of our moving our eyes many times per second, information from the retinas comes in the form of rapid snapshots of different parts of our environment.⁵⁷ Not only must the snapshots be pieced together to construct the visual scene, but the noisy retinal inputs resulting from our saccades must be correctly interpreted as being due to our own eye movements rather than changes in the environment.

That our experience is not herky-jerky despite these eye movements indicates what kinds of processing takes place between light hitting our retinas and our having experience. Causal interactions between our sense organs and the external world yield sensory input, which initiate a series of stages of processing, resulting in perceptual experience. Considering saccades helps us understand what sort of processing is involved, adding layers to our scientific causal picture. Visual processing is not akin to an assembly line whereby a continuous stream of information about the visual scene from both eyes in two dimensions arrives on the conveyor belt and is reconstructed in a straightforward step-by-step process, resulting in our uninterrupted three-dimensional experience of the world. When we saccade, information from the eyes is suppressed so that input does not enter into awareness. This is why we cannot detect our own eye movements when looking in a mirror. Suppression—which begins just prior to a saccade—

⁵⁷ Due to the distribution of photoreceptors on the retina these snapshots are not homogenous in terms of resolution. This further complicates the processing involved in producing our visual experience and may even introduce certain puzzles regarding consciousness, however these issues are beyond the scope of my concerns.

makes sense, as it prevents us from experiencing motion and blurriness each time we saccade. However, suppressing information from the eyes during saccades means that there will be brief gaps in the stream of incoming information to the visual system. Because of the time delay from light hitting our retinas to visual experience and the speed with which we saccade, the brain has a chance to construct visual experience from the incomplete parts. Hence, though the brain is effective in producing a stable and continuous visual experience, the process is far from straightforward.

The brain may be able to faithfully reproduce the external world based on the inputs it receives from the eyes. By itself, the sort of construction necessitated by saccades does not imply that there will be mismatches, that there will be discrepancies between our experience and the external world. But once we move away from the simple conveyor belt and discover how our visual experience is produced, we can see how it is possible for mismatches to result from the sorts of processes being considered here. As the brain stitches snapshots together into the spatiotemporal fabric of experience, there will be some inputs missing due to saccading. However, we do not have blank periods in our visual experience because the brain fills in the temporal gaps during which input is suppressed. If we are viewing a static scene, then lacking information from brief periods may not make a difference. But temporal filling-in when we are viewing motion is another matter.

Evidence of temporal filling-in comes from studies like that conducted by Yarrow et al. (2001). Here, subjects saccade to a digital counter that begins to count from zero with the initiation of their eye movement. Subjects overestimate the duration during which the first number is displayed, as compared with subsequent numbers. The amount of time subjects overestimate seeing the number roughly corresponds to the duration of their saccade.

The chronostasis observed in the study is the same sort of illusion that occurs when we look towards a clock and the second hand appears to initially stick, or take longer than usual to move. As the authors note,

...the illusion occurs because the brain extends the percept of the saccadic target backwards in time to just before the onset of the saccade. ...We suggest that temporal extension of the target's percept is one of the mechanisms that 'fill in' the perceptual 'gap' during saccadic suppression. (p. 302)

Identifying and measuring the effect is made possible—or at least made easier—by the metronomic movement of a clock displaying seconds. If we saccade towards something that does not move in a uniform way we likely will not notice that something has been briefly stretched in time; similarly, if the motion is too slow, as with a minute hand, the duration added by filling in the gap resulting from the brief saccade will be insignificant compared to how long the hand is in place. Thus, as the authors say, the extension of the percept backwards in time “occurs every time we move the eyes but is only perceived when an external time reference alerts us to the phenomenon” (p. 302).⁵⁸

Whether or not a situation affords us the opportunity to notice temporal filling-in, such a process occurs. That is not to say that temporal filling-in always results in a mismatch. If I saccade to my stationary fridge and my experience of the fridge is extended backwards in time so that I do not notice suppression of input during my saccade, there will be no mismatch.⁵⁹ But in

⁵⁸ Perhaps “perceived” is not the best word here, and “noticed” would be more appropriate.

⁵⁹ Some may consider experience to be in a sense illusory because we do not notice the gap of the saccade, and such phenomena prompt some to note that conscious experience is not what it seems or that experience is constructed. But I am not arguing that we must invoke mental representations to explain the seamlessness of our experience. The lack of temporal gaps in our experience helps shed light on perceptual processing, but this concerns the relationship between

the Yarrow et al. study there is a mismatch. The post-saccadic input is stretched backwards in time to fill in the duration of the saccade, when input was suppressed. As a result, in our experience the number (or the second hand) is there longer than it actually was—that is, in our experience it is somewhere that, for a brief period, it was not. The brain can extend the percept backwards in time because of the time delay inherent in perceptual processing. It takes time for our brains to process sensory input, and our experience lags slightly behind what is going on in the world. This lag makes it possible for the brain to “extend the percept backwards in time.” If we were to plot out a timeline of what is going on in the world next to a timeline of our visual experience and discounted the time delay due to our perceptual processing such that the two lined up, then the mismatch would be apparent—at a certain point in the sequence of our experience the number will already have changed, though that had yet to happen in the external world.⁶⁰

Similarly, when things are moving around us in daily life our visual experience of them can be slightly distorted due to temporal filling-in. Because the time periods being filled in are so brief such mismatches will be minor and inconsequential. And because we are not typically looking towards clocks the mismatches are sure to go unnoticed. But that does not mean such mismatches are uncommon, for our world is far from static.

sensory inputs and experience. Mismatches concern the relationship between our experience and external objects and events.

⁶⁰ The point being made here is not that a mismatch results simply because there is a delay between events in the world and our experience of those events, thus we experience things being in certain locations they no longer are in (or even possibly experience things that no longer exist, such as faraway stars). Just as we can highlight a mismatch by comparing what some experience is like in the spatial dimension and what the external world is like spatially, so we can compare two timelines to make salient the mismatch. With temporal filling-in, mismatches result when our experience differs from what is external to us—e.g. the actual motions of external things.

2.4. Visual Completion

Visual completion includes a variety of related phenomena ranging from filling-in of the blind spots and scotomata, to illusory contours and the filling-in of illusory surfaces and shapes, to different effects resulting from the extended fixation of a stimulus.⁶¹

Despite the lack of photoreceptors where the optic nerve exits the retina, we do not experience a gap in our visual field. This is largely because the two blind spots do not overlap and in normal binocular vision information from the opposite eye will compensate for the lack of input from the blind spot. But even if we explore our environment with one eye closed there is no gap in our visual field because the blind spot is filled in. And even in monocular vision this filling-in is likely to go unnoticed.

We can detect filling-in of the blind spot when certain stimuli are viewed under particular conditions. For example, a centimeter-wide black spot on a white sheet of paper will disappear if viewed at a distance of about a foot, displaced slightly towards the side of the open eye. In order to study this phenomenon, a variety of stimuli have been created which render the filling-in salient (for examples, see Ramachandran, 1992; Komatsu, 2006). Monocular viewing of these stimuli at a certain angle and distance results in the surrounding colour or pattern filling-in the dissimilar portion when it falls in the blind portion of the visual field. Thus, when viewing a red square with a blue circle at its centre, one experiences only a red square.

Similarly instantaneous and stimulus-independent filling-in also occurs with scotomata, or portions of the visual field where information is lacking due to damage to the visual system.

⁶¹ See Pessoa et al. (1998), Komatsu (2006), and Weil and Rees (2011) for reviews and discussion of such phenomena and various proposed mechanisms, as well as different taxonomies.

Once again, the brain fills in these regions with information from the surround and there is no gap in one's experience.⁶²

The brain areas involved in filling-in of the blind spot and of scotomata have been studied extensively. Such filling-in involves active processes, whereby the filled in portion of the percept is actively constructed by the brain (Komatsu, 2006).⁶³ That the brain uses information from the surround in order to complete the filling-in and that patterns in the surround can be continued through the filled in portion of our experience suggests that the brain may try to construct what is most likely to be external to us, based on its input.

While there is a lot of interest in this sort of filling-in, in what ways such findings shed light on normal visual processing and everyday perceptual experience is far from clear. Studies investigating filling-in at the blind spot typically involve mismatches. When a subject's visual experience differs from the stimulus before them, experimenters can be sure that what appears in the portion of the visual field corresponding to the blind spot is filled in by the brain. Though we may only notice how our visual experience differs from what is external to us when we view particular stimuli, monocular viewing will typically result in mismatches at the blind spot, where

⁶² A similar phenomenon has been observed in auditory perception, including the perception of speech. When a sound is interrupted by noise—or when a brief segment of a signal is artificially replaced by a segment of noise—we experience the sound as continuing uninterrupted (Petkov et al., 2007; Shahin et al., 2009). Studying this phenomenon helps us understand everyday auditory perception in environments with background noise and interference.

⁶³ Dennett (1992) raises various concerns about what he takes to be the widespread use of “filling-in” in the scientific literature. His discussion covers a variety of possibly unrelated phenomena, including filling-in at the blind spot. Though he may be right to question the use of such talk as it relates to certain phenomena, and raises important questions about consciousness—e.g. relating to peripheral vision—recent empirical work supports the claim that there is filling-in of the blind spot by the brain. In particular, when the blind spot is filled in in monocular viewing—including when this results in a mismatch—there is activity in the specific area of V1 corresponding to the blind portion of the visual field, where there is no input (Komatsu, 2006). (For a reply to Dennett, see Churchland and Ramachandran, 1994.)

filling-in occurs.⁶⁴ However, binocular viewing is the norm, as we typically do not go around with one eye closed. It seems reasonable to think that filling-in occurs even when both eyes are open, but it is difficult to say. Given that the early visual areas of each brain hemisphere receive inputs from both eyes, it may be difficult to determine whether filling-in occurs in such situations, since the area of the visual field corresponding to the blind spot will have input due to the information from the other eye. Simply by looking at, say, V1, how would we determine if the information encoding the relevant portion of the contralateral visual field was filled-in or due to the other eye? But regardless of whether filling-in at the blind spot and of scotomata occurs when both eyes are open, since there is information from the opposite eye available, this sort of filling-in will not result in mismatches in our daily lives.

Illusory contours are edges, lines, forms, and volumes we experience in the absence of physical support or visual information, when there is no correlate in the physical world. Like filling-in at the blind spot, the effect is immediate and does not require extended fixation; however, unlike such filling-in, the perception of illusory contours only occurs with certain stimuli. Kanizsa figures induce illusory contours and the experience of a shape formed by such lines. For example, four black pacman shapes arranged appropriately on a white background induce the experience of a white square. Lines and gratings arranged appropriately induce similar effects (for examples, see Weil and Rees, 2011; Murray and Herrmann, 2013). Though stimuli

⁶⁴ As noted, if we look around our environment with one eye closed we will not notice that our blind spot is being filled in or that our experience may therefore differ in subtle ways from the external world. Monocular viewing need not always lead to a mismatch, as what is filled in may correspond to what is in the portion of our visual field—for example, when staring at a uniformly-coloured wall. However, even when the brain gets it right—so to speak—the lack of causal connection between a portion of the scene and the corresponding part of our experience would seem to pose a problem for the naïve realist analysis of perception. The sort of situation described here may be akin to so-called “veridical hallucinations” (Macpherson, 2013).

are typically created in order to induce and demonstrate the most robust effects, a wide range of stimuli will lead to the perception of illusory contours.

A related effect is the spreading of brightness or colour through a shape or surface formed by illusory contours. For example, neon colour spreading will occur with a stimulus similar to the Kanizsa figure described above, where the relevant portion of each circle is neon-coloured rather than empty. Here, the colour spreads to fill the illusory square, which has a neon tinge. The effect disappears if two adjacent inducers are covered.

Like filling-in at the blind spot, such effects involve active processes and feedback to early visual areas, as well as activity in early visual areas corresponding to the filled in lines and shapes (Komatsu, 2006; Weil and Rees, 2011; Wokke et al., 2013; Murray and Herrmann, 2013). It is thought that the processes involved in these effects are also involved in the normal perception of objects and surfaces, which typically occurs in complex environments where objects are occluded in various ways. The sorts of figures used to clearly induce these effects exploit mechanisms that are involved in everyday object identification and recognition, which more and more are taken to involve active neural processes and feedback from later to earlier stages of processing.

Although the processes underlying illusory contours may be involved in normal perceptual processing in the brain, it does not follow that in normal conditions illusory contours are ever-present. Nevertheless, unlike the sorts of mismatches we see with filling-in at the blind spot, it does seem like mismatches due to illusory contours will occur regularly in everyday experience. Normal environments include many objects at various distances and angles from us, layered in complex ways. As a result, such complex scenes are likely to involve lines, boundaries, and arrangements that induce illusory contours. The illusory contours we experience

in everyday settings may not be as pronounced as the neon colour spreading seen with powerful inducers, and we will not notice the lines were illusory as we move about in our environment and our point of view shifts and we receive new visual input. But it is reasonable to expect that with certain neural mechanisms and tendencies present and the complex input the brain receives, there will be boundaries, lines, and possibly even volumes we briefly experience when none are present in the external world.

Lastly, there are various visual effects that result from steady, extended fixation of a stimulus for several seconds, like image stabilization and the uniformity illusion (Komatsu, 2006; Otten et al., 2017). Given that these effects depend on such extended fixation it can be unclear how they relate to everyday perception, and we certainly cannot expect them to result in mismatches in daily life.

2.5. McLaughlin on Normal Geometrical Misperception

McLaughlin (2016) argues that “normal visual misperception is systematic and widespread” (p. 235). By “normal visual misperception” he means visual illusions that occur in normal viewing conditions, where a visual illusion occurs when there is a difference between our visual experience of some thing and how that thing actually is.

When we have a visual illusion in virtue of something looking some way
W to us, the something in question is not W, and so we have visually
misperceived it. (p. 233)

McLaughlin’s visual illusions or cases of normal misperception are therefore mismatches—they are instances where our experience differs from what is external to us, as in when we experience some object having a property it lacks. He argues that such mismatches are not exceptional and that

...the visual errors in question don't occur only in, or only mainly in, ecologically aberrant circumstances. They are commonplace in our ecological niche. (p. 235)

McLaughlin discusses a variety of findings relating to visual illusions. Some of these findings concern illusions discovered over a century ago, while others involve phenomena that have been investigated only recently. Vertical and horizontal lines of equal length typically do not look equal to us, with the vertical ones appearing to be longer. This is evident in the upside-down T-illusion, where a vertical line bisects a horizontal line of the same length, as well as when lines of equal length are arranged as a fan. Relatedly, squares and circles look slightly stretched to us, elongated vertically. Angles also appear slightly distorted—obtuse angles appearing larger than they actually are, acute ones smaller—in particular within certain contexts, as demonstrated with figures like the Zöllner illusion. In the Pogendorff illusion a straight line passing behind one or more occluders will not look straight, making it possible to mistake a segment as belonging to a nearby line that stops behind an occluder. Meanwhile, when objects are partly occluded the visible portion will look larger than it actually is (and larger than the portion looks when viewed in isolation or when not partly occluded). Finally, he notes how in foveal vision there are various distortions of space, such as things appearing to be farther apart than they actually are (whereas in peripheral vision they will appear closer together than they are).

Crucially, both for our purposes and for him to make his case, the illusions McLaughlin points to involve stimuli that resemble arrangements we regularly encounter in our daily lives. The figures used to illustrate the various visual illusions resemble things we find in natural scenes. For example, the size illusion caused by occlusion mimics the partial occlusion of objects that is omnipresent in everyday perception. Such normal misperception results from the normal

workings of our visual system and how it processes visual information.⁶⁵ As a result of this processing, we can expect these sorts of mismatches to occur in day-to-day life.

Despite the prevalence of the sort of misperception McLaughlin draws attention to, it would be wrong to think our visual experience is largely mistaken or illusory. As he notes:

The normal visual geometric misperceptions that I've highlighted don't pose a practical problem. There's much that vision gets right. Needless to say, relying on vision, we do just fine behaviourally, and we do spectacularly well cognitively. Vision science, a cognitive achievement, has even shown us the error of our visual ways. (p. 291)

Though such misperception will certainly result in mismatches in our day-to-day lives, for the most part these will be relatively minor. They will include cases where something appears slightly larger than it is, or where we mistake a part of some distal object as belonging to a different object, or where our experience of an object's location is slightly off. Then, if we view things from a different position, or see an object when it is not occluded in a certain way, our experience will be different. And, typically, the differences between our successive experiences will be slight and go unnoticed.

⁶⁵ Roughly, McLaughlin understands our visual experience as being the solution to the inverse optics problem. This is the problem of generating visual experience on the basis of the two-dimensional retinal images induced by objects and events in the world. The problem arises because a retinal image includes information about shape, angle, orientation, and other things; and, there are infinitely many possible scenes that would lead to the same retinal image. The processing of information in the visual system takes us from images on our retinas to visual experience. The degree to which we “solve” the inverse optics problem may be seen as the degree to which our experience of the external world is accurate. Obviously we do fairly well and get by fine, in particular given how fast the processing is and needs to be. But that does not mean we get things right, much less all of the time. McLaughlin discusses a few different theories about how we solve the problem, and insofar as they posit different principles underlying our perceptual processing, these different theories will offer different explanations of normal geometrical misperception.

2.6. Colour Perception

Recent work on colour perception reveals how our experience of colour depends on a variety of factors besides the properties of the physical objects we perceive. Most fundamentally, “the light reflected from an object’s surface entangles information about reflectance with the information about the incident illumination” (Radonjic et al., 2015, p. 1).⁶⁶ Disentangling these two factors is essential in order to achieve colour constancy, i.e. the stability of an object’s perceived colour across environments. But the illumination of the scene is not the only factor that can impact our colour experience. The surrounding objects and surfaces also influence our experience of a target stimuli (for example, see Allred and Olkkonen, 2013). Meanwhile, environment-independent factors like memory and previously viewed stimuli can impact colour perception (Olkkonen and Allred, 2014; Olkkonen et al., 2014; Olkkonen et al., 2016).

Consider contrast studies, where an object or patch with certain reflectance properties is placed alongside or surrounded by different stimuli. In the basic simultaneous colour contrast case, two identical patches will be surrounded by different borders and placed next to each other. As a result of the different surrounds, the identical central patches will look different to us.⁶⁷ The same effect is achieved with the variety of other stimuli that have been produced to study such phenomena, but the basic case is enough to demonstrate that our experience of an object with certain surface reflectance will differ depending on the surrounding objects and environment. The converse also happens, as demonstrated by studies using the basic paradigm: two objects with different reflectance properties can look to be the same colour due to their surrounds.

⁶⁶ This creates an inverse optics problem, of the sort discussed in the previous footnote.

⁶⁷ The same design can be used with grayscale patches of varying lightness to similarly influence apparent lightness.

This empirical work shows that an object that typically causes a certain colour experience can, in a different context, cause a very different colour experience; and, two objects that typically cause different colour experiences can, in certain contexts, cause identical colour experiences. Though colour constancy is often achieved, such studies reveal that there are often failures in constancy. As noted, it is not only the surrounding objects that influence colour perception, but illumination and perceiver-dependent factors like perspective and memory. And while the stimuli used in certain studies may be different from real-world objects and scenes, similar sorts of conditions are present in daily life. In many ways, real-world scenes and objects introduce many more complexities than are present in experimental settings (Brainard and Radonjic, 2014).

Now, does this mean that colour perception involves mismatches? Because there is significant debate about how we should understand colour, I am hesitant to claim that when it comes to colour our experience often differs from what is external to us. Doing so would seem to involve taking a stance on what exactly is “out there” in the case of colour. However, regardless of how we understand the metaphysics of colour, I think we can view the sort of variation we see in the case of colour perception as being, for our purposes, analogous to the mismatches discussed previously.

With mismatches there is a difference between what a person experiences and what is external to them. In the flash and beeps case we experience two flashes even though there is only one. With the McGurk effect we hear [da] when lip movements for [ga] are accompanied with audio of [ba]. In such cases, objects appear to us to have certain properties they actually lack or we have the experience of things that are not actually out there in the world. The issue—as far as we are concerned—is how we are to explain such cases, and whether we can explain what we

experience simply in terms of our being aware of the external objects and their properties. The variability we see with colour raises the same issue regarding explanation. Our colour experience can vary even if the target stimuli remains the same. Meanwhile, objects with different surface reflectance can cause identical colour experiences. Thus, it seems like we cannot simply say that our colour experiences involve the awareness of mind-independent objects and their properties; we will have trouble explaining these experiences simply in terms of the relevant external objects. Because colour experience depends on a variety of factors—some external to the subject, some internal—that contribute to a lack of constancy, we can get the same sort of discrepancy between external objects and their properties and our experience as we do with mismatches.⁶⁸

2.7. Speech Perception

Studies investigating the perception of speech reveal something similar to what I remarked on regarding colour perception. With speech, there is a lack of correspondence between the acoustic signals in the external world and our resulting auditory experiences.

Appelbaum (1998, 1999) discusses a variety of findings comparing acoustic speech signals and speech perception, typically focusing on phonetic segments. The most significant findings are that different acoustic signals and cues can lead to the same auditory experience, whereas a single acoustic pattern can cause different auditory experiences in different contexts.

There are a number of factors that have been shown to influence our perception of speech. If a

⁶⁸ To borrow language from Pautz—as seen in the previous chapter—we may say that there is poor external correlation, but good internal correlation, since our experiences depend on our neural processing. However, unlike his correlations between similarity relations, the poor external correlation here calls into question the naïve realist’s ability to explain the character of a given experience in terms of an object’s properties because the relationship between the properties of objects and our corresponding colour experiences comes apart.

sound is ambiguous between two phonemes, only one of which will complete a word, we tend to hear the signal as the phoneme that forms the word. Meanwhile, identical acoustic material can be heard as different word portions depending on where it appears within the stream of a sentence.⁶⁹ The same acoustic information may be perceived as different vowel sounds depending on the speaker that produced the output and how they vocalize other sounds. Semantic information and noise in the environment can also shape our perception of incoming speech signals, and whether or not a signal is experienced as speech can depend on expectation. Of note, many of these effects depend on an individual being a speaker of the relevant language, and will therefore vary accordingly across individuals.

Thus, there is the same sort of variability between physical stimuli and our experience as we saw above. Our auditory experience can vary even if the acoustic stimulus remains the same, and different acoustic stimuli can cause identical experiences. The issue is not the complexity of the processing involved in speech perception, but rather the relationship between our perceptual experience and the external objects and events that cause our experience. As with colour, this variability calls into question whether we can explain what we experience simply in terms of our being aware of external objects and events.

3. The Argument that Mismatches are Ubiquitous

The mismatch argument begins with a mismatch and attempts to generalize our explanation of that case to all perceptual experience. But, as we saw in the previous section,

⁶⁹ This effect is due to the syntactic structure of a sentence, which itself is not straightforwardly determined by the acoustic signal—for instance, the syntactic structure will determine whether a pause of a certain duration is interpreted as a break between clauses or merely between words.

mismatches are not rare occurrences; rather, they are ubiquitous.⁷⁰ This would seem to eliminate the need to generalize from one case to many, suggesting another way to start with a mismatch and argue for the same conclusion.

Here is that second argument:

(i) In one example of the McGurk effect, an individual is presented with lip movements for [ga] and audio of [ba]. These are the mind-independent stimuli causally interacting with the individual's sense-organs.

(ii) The individual experiences a different syllable altogether, [da]. Their auditory experience is of the syllable [da], and they do not experience either [ba] or [ga].

(iii) We can explain what the individual experiences if we suppose that the brain constructs a representation that determines the character of their experience. The individual is aware of this representation. We cannot explain the character of the individual's experience simply by pointing to the external objects, since their experience, or what appears to them, differs in the way it does from the external objects. In this case, the individual is not immediately aware of the external objects.

(iv*) Such mismatches between the objects and events external to the individual and the individual's experience are the norm. If the character of our experience in these instances is determined by a representation constructed by the brain, then the character of our experience is normally determined by a representation constructed by the brain.

⁷⁰ Most of the empirical work discussed concerned vision (though the discussion of speech perception related to audition and multimodal perception concerns all of the modalities). This is unsurprising given how much work is done on vision as compared to the other modalities. It seems likely that as we continue to investigate audition, touch, taste, and smell we will find evidence of further mismatches occurring in these modalities. However, in what follows I rely only on the work discussed in the previous section.

(v) Therefore, in perceptual experience we are aware of a mental representation generated by the brain; we are never aware of mind-independent objects and their properties.

The conclusion is the same as in the mismatch argument, and from this conclusion it follows that naïve realism—which says that perception involves the direct awareness of mind-independent objects, which determine the character of our experience—is false. Rather, perception is a matter of mentally representing the world, and the character of our experience is determined by the representation. Once again, indirect realism is vindicated.⁷¹

As with the mismatch argument, we can start the argument that mismatches are ubiquitous with any mismatch. In the first two premises above I describe the McGurk effect, but any mismatch will do.

The two arguments also share (iii). As we have seen, the issue here is how to best explain our experiences in such cases. Whether we should accept (iii) depends on whether or not we are justified in explaining these experiences with mental representations.

The difference between the arguments is the fourth premise—instead of (iv) we have (iv*). Our examination of recent empirical work in the previous section supports (iv*). By looking at the different situations in which, and the different ways in which, mismatches result, we see that mismatches are ubiquitous. Whether or not we should accept (iv*) therefore seems like a straightforward empirical question, and the evidence supports its acceptance. However, the naïve realist may concede the point, yet still try to respond to (iv*). It also seems like we may

⁷¹ As noted in the previous chapter, what we are aware of in perception is what determines the character of our experience. In (v), we may want to say that we are never *directly* aware of mind-independent objects, instead being directly aware of representations generated by the brain. Opting for such wording changes nothing, as what we are directly or immediately aware of would be what determines the character of our experience.

need to say a bit more in order to move from (iv*) to (v). These are two things I will consider below.

We now appear to have two compelling arguments that marshal recent empirical work against naïve realism and in support of indirect realism. In the next three sections I will scrutinize the three key claims in these arguments. I begin by looking at (iii) and whether we are justified in invoking mental representations to explain these experiences. I then consider (iv) and ask whether we are right to think that perceptual processing always operates in the same way. Finally, I examine (iv*) and consider various ways the naïve realist may respond given that mismatches are ubiquitous.

4. Explanation and Representation

In certain instances there is a mismatch between the objects and events external to an individual and the experience of the individual caused by those external objects and events. For example, with the McGurk effect an individual may be presented with video of a person saying [ga] synced with audio of [ba], and their auditory experience will be of the syllable [da]. In the flash and beeps study, there is a single flash of light, but the individual experiences, visually, two flashes. In such cases there is a difference between what is external to the individual and what they experience, between the properties of the relevant mind-independent objects and events and the properties that they are aware of in perception, the sensory qualities that appear to them. The question we will be concerned with in this section is: how can we best explain these experiences?

According to (iii), we should explain our experience in such cases by claiming there is a mental representation that determines the character of these experiences. The explanation is that

the brain constructs a representation, and the qualities that appear to us in experience are determined by, or due to, that representation.

A possible alternative explanation would be that we are directly aware of external objects and events, and their properties, and our experience in such cases is the way it is because of this awareness. But given that our experience differs in the way it does from the external objects and events, it seems like this account will not do. This alternative does not explain how or why our experience has the phenomenal character it actually does in such cases. The possible alternative, then, does not offer a viable explanation.

Since pointing to external objects and events will not do, we explain such experiences in terms of things internal to the individual. Though an account can be offered that is focused solely on the mind and how the individual is representing the world in such cases, it is only natural that the scientifically-inclined theorist will focus on the role of the brain in perception. Given what we know about the brain, the account offered in (iii) is the best explanation of our experience in such cases.

The relationship between the brain and experience is a tight one. We know that different patterns of brain activity reflect different experiences. Different brain areas subserve different functions and different sorts of experiences, and recordings of individual cells demonstrate on a fine-grained level how activity in small neural populations is correlated with the experience of different properties. Consider how the area V1 is organized retinotopically, meaning that activity in different areas of V1 is caused by stimuli in different parts of the visual field and also corresponds to our experiencing something in particular parts of our visual field. When the blind spot is filled in in monocular viewing, there is activity in the specific area of V1 corresponding to the blind portion of the visual field, where we experience what is filled in (Komatsu, 2006).

Since our experiences are reflected in our brain activity, and the sensory qualities that appear to us correspond to particular brain activity, we can understand this as our brain representing the world or the scene, such that the activity of different groups of neurons is associated with different aspects of the scene and different phenomenal properties. What we experience—the properties we experience—depends on what is in the brain, on this brain activity. That is to say, our experience is determined by the representation that is constructed—or generated—via the processing of sensory input. The culmination of perceptual processing is a representation of the scene based on the input, and our experience is determined by this mental representation.⁷²

The naïve realist may be happy to accept (iii). Even though they think that perceptual experience involves the direct awareness of mind-independent objects, they may be willing to concede that in cases where there is a mismatch between what an individual experiences and what is external to them, that in these cases the individual is directly aware of a mental representation rather than mind-independent objects. If they take mismatches to be illusions and do not view illusions as cases of perception, then they will not be troubled by our explanation of such cases.

Martin, as we saw in the first chapter, groups illusions and hallucinations. He adopts a disjunctivism that distinguishes between veridical perception on the one hand, and illusions and hallucinations on the other. Being a disjunctivist, the account he offers for cases of perception does not apply to the bad cases. And he is free to accept any account of illusions and of hallucinations, even one invoking mental representations in the way we see above. Therefore, a

⁷² By “culmination” here I do not necessarily mean the end point in time, nor even the last stage of processing the input may reach in the brain if we think of the input as undergoing successive stages of processing—though, both may very well be true of the representation. Rather, given its relationship to our perceptual experience and the role the representation plays in our psychological life, it makes sense to think of it as the result of perceptual processing.

naïve realist who, like Martin, can accept that in cases of illusion we are not directly aware of mind-independent objects may not care what we say about mismatches.⁷³

But a naïve realist may not want to accept (iii). After all, not all naïve realists group illusions and hallucinations—some classify illusory cases alongside cases of veridical perception. Such a theorist is going to reject an explanation of mismatches that invokes mental representations. And, regardless of how they classify illusions, given the subsequent premises in our two arguments, a naïve realist may think their best chance for confronting the arguments is to dispute (iii).

If someone wants to claim that we are not justified in invoking mental representations to explain our experience in mismatch cases, then they must provide an adequate explanation of these cases that does not feature mental representations. If we want to see whether a naïve realist can provide such an explanation, we must look to naïve realist accounts of illusions.⁷⁴

We may think the naïve realist is in trouble here, considering that a number of people argue that naïve realists cannot offer a compelling explanation of illusory experiences that is consistent with the core commitments of their theory (see, for example, Brown, 2008; McLaughlin, 2010; Millar, 2015). The naïve realist's alleged inability to accommodate or account for illusions is touted as a reason to reject the theory; the idea being that we should

⁷³ Below, in section 6, I will consider the question of whether we should take mismatches to be cases of illusion. I am not endorsing such a designation here. However, I think it is clear that people will typically consider the McGurk effect and the flash and beeps case illusions (if not hallucinations). And given how I have characterized mismatches and how most people define illusions, it seems reasonable to think most people will consider mismatches in general to be illusions. Where naïve realists may differ is on whether illusions are grouped with hallucinations or cases of veridical perception.

⁷⁴ Once again, I am not claiming that we should consider mismatches to be illusions. Rather, the point is that if we want to see how naïve realists will try to explain our experience in cases of mismatches we should look at their accounts of illusions, which effectively are attempts to extend the naïve realist account of perception to cases where there is a difference between what an individual experiences and what is external to them.

prefer theories of perception that can account for all of the phenomena pertaining to experience over ones that cannot. Though it seems reasonable to want a theory of perception to have something to say about the Müller-Lyer illusion, I am not criticizing naïve realism for any apparent lacuna regarding illusions. Here, we are interested in accounts of illusions offered by naïve realists in order to see if any such account can help them dispute (iii). It is in this spirit that I look to these accounts of illusions, to see how naïve realists may try to explain mismatches.

Fish (2009) divides illusions into three main classes and offers a different account for each.⁷⁵ “Physical illusions” include cases where objects will look to be of different shapes or colours when viewed in particular conditions, as from certain perspectives, with certain illumination, or through things like water or lenses. Such illusions are due to how things in the world are, external to the subject, and Fish notes that in order to explain why things appear to us in the manner they do we need only point to physical things in the environment.

The reason that there is an illusory appearance in the case of physical illusions, then, is purely a matter of what is going on in the world and how that affects the patterns of light that impinge upon the subject. (p. 148)

Such illusions are predictable in the sense that if we know the physical arrangement of the scene and the position of the subject we will be able to know that they will have an illusory experience and how things will appear to them.

An example of a physical illusion he discusses in detail is looking at a circular object through a lens, such that it looks elliptical. The experience of looking at the object through the lens is different from the experience of looking at the object without the lens—the experiences

⁷⁵ He focuses only on vision, though it is not difficult to see how his accounts would generalize to the other modalities. These categories are meant to be rough, and he notes that for certain cases it will be unclear which category it belongs in.

have different phenomenal characters. Fish explains the difference in character by pointing to the different properties we are aware of in the two cases.⁷⁶ When we look at the circular object we are aware of various properties including its shape; however, when we look at the object through the lens we are aware of various properties such as its colour, but we are not aware of its shape. Fish compares the lens to a microscope, which can make us aware of certain properties while simultaneously preventing us from being aware of other properties of an object that we may normally be aware of. When we look at the object through the lens we cannot be aware of the object's elliptical shape, as it does not have an elliptical shape. According to Fish, we are not aware of the object's shape at all and therefore our experience has no phenomenal properties corresponding to the object's shape.⁷⁷ We might think that we have the experience of an elliptical shape, but this is mistaken. Rather, we merely judge that the object is elliptical.

A second class of illusions are “cognitive illusions.” These are cases where “we see something...but take it to be something it is not” (p. 149). In contrast to physical illusions, cognitive illusions are due to how things are with the subject. Fish discusses the example of seeing a rope but taking it to be a snake—we can imagine a person who has a fear of snakes walks into a cave and spots a coil of rope under a log. In such a case, the subject is aware of the rope and its various properties, but fails to perceive that it is a rope. The subject takes themselves to see a snake because—primed by their fear of snakes, the dimness of the cave, the way the rope

⁷⁶ Fish often puts things in terms of being acquainted with certain facts rather than being aware of certain properties, but the difference is unimportant and for ease of discussion I will stick with the latter. On his view, an experience has a phenomenal character or, alternatively, an array of phenomenal properties, “each of which involves...acquaintance to a worldly fact” (p. 148).

⁷⁷ In Fish's terminology, the distorting lens “prevents the subject being acquainted with facts concerning the shapes of objects seen through the lens” (p. 164).

is coiled, and other factors—they mistakenly take the rope to be a snake.⁷⁸ Their taking themselves to see a snake is a purely cognitive matter. Fish notes that the character of their experience would be no different if they merely failed to recognize that the object was a rope.

“Optical illusions” are somewhere in the middle of the previous two, though closer to physical illusions. Like physical illusions, optical illusions depend largely on how things are in the external world and are relatively predictable. However, these illusions are also due in part to our particular perceptual processes, and in order to explain these illusions we will have to look at what is going on inside the subject. For example, Fish discusses the Müller-Lyer illusion. Here, the lines look to be of different lengths, which cannot be due to our awareness of some properties of the lines because the lines are actually the same length. But the lines continue to look to be of different lengths even though we know they are the same length—this is what distinguishes such illusions from cognitive illusions.⁷⁹ What is happening here, says Fish, is that we instead believe that it looks like one line is longer than the other; this belief explains why we claim that the two lines look to be different. But this belief does not make it such that the phenomenal character of our experience as it relates only to the horizontal lines is the same as it would be if we were

⁷⁸ More specifically, Fish argues that they (mis)deploy their snake concept, or their snake “conceptual-recognitional capacity.” This explains why they fail to see the object as a rope—they deploy the wrong concept—and why they falsely believe there is a snake. Importantly, however, the deployment of concepts does not by itself contribute to an experience’s phenomenal character. It may allow us to become aware of some property of an object, but not become aware of phenomenal properties that the objects we perceive do not have (p. 169-72).

⁷⁹ The difference with physical illusions can be seen by the fact that the retinal images produced by the two horizontal lines will be the same, thus we “cannot explain our susceptibility to the illusion purely by appeal to how things in the world affect the light impinging upon our retinas” (p. 173). But while we may understand Fish distinguishing between physical and optical illusions on this basis, the distorted lens case and the Müller-Lyer illusion may not seem that different to us if we are concerned with how things appear to us in the cases. Since I am unsure I would consider a circle seen through a distorted lens to be an illusion—or a stick in water, another of Fish’s examples—I hesitate to draw similarities between physical and optical illusions. But I will note that even in the case of physical illusions Fish allows for our judgment—something internal to the subject—to sometimes play a role.

looking at horizontal lines of unequal length. As when a rope is mistaken for a snake, the false belief does not alter the phenomenal character of the experience. This is evident based on what he says about the Kanizsa triangle. Here, the phenomenal features of our experience concern only the incomplete circles.

The illusion is that, rather than seeing incomplete circles, we “see” a triangle of the same color as the background partially occluding completed circles. What it is like to have an experience of this figure is, I suggest, a matter of the subject’s having an experience with an array of phenomenal properties corresponding to the range of facts concerning the incomplete circles. However, because we are primed to see the world in terms of bounded objects, these features of the figure have the consequence that the subject’s triangle conceptual-recognitional capacity is erroneously deployed, which in turn explains why the subject falsely judges that there is, or at least that it looks as though there is, a white triangle occluding the rest of the scene. And because the misleading aspects of the figure operate at such a low level, their contribution cannot be offset by higher cognitive factors. (p. 177)

I will now look at whether Fish’s view of illusions can help us explain mismatches.

It seems clear that most—if not all—mismatches will be optical illusions. The various mismatches between our experience and what is external to us result from processing that takes place once we have received sensory input, which means they will not be physical illusions. And unlike cognitive illusions, mismatches will persist even if we fully understand what is going on in the given situation—we cannot become impervious to the McGurk effect by dropping some erroneous belief. The sorts of arrangements that lead to mismatches may be much more complex and diverse than the ones Fish considers, and due to individual differences in experience and the

development of perceptual systems it is possible to have some variation between subjects. But, as with optical illusions, mismatches depend on the interaction between how things are in the external world and our processing of the resulting sensory input.

One way of understanding Fish's account of optical illusions is that when stimuli are arranged in a certain way they affect our perceptual processing in certain ways, such that the mechanisms involved in our perceptual processing prevent us from becoming aware of certain properties of the scene. The effect of our perceptual processing here is analogous to that of the lens, which prevents us from being aware of the object's shape. In the case of the Müller-Lyer illusion, certain features of the stimulus impact our perceptual processing in such a way that we are prevented from being aware of the length of the lines.⁸⁰ This would explain why we do not experience the patchwork images, or why we do not hear [ba], or why we are not aware of the blue circle at the center of the red square in monocular viewing when it falls in our blind spot. In these situations our perceptual systems are exploited in certain ways, such that we are prevented from becoming aware of certain properties.

But what about the relevant properties we do experience in these cases? In the binocular rivalry study our experience is of alternating coherent images, in the McGurk case we experience [da], and with filling-in we experience red in the blind portion of the visual field. According to Fish's view, such phenomenal properties are not actually present in these experiences. Rather than offering an explanation of the character of our experience when it differs from what is

⁸⁰ Millar (2015) interprets Fish in this way. We also may think certain perceptual mechanisms would allow us to become aware of properties we would not otherwise be aware of, as is the case with the microscope. If we are to understand (some) mismatches in this way, that would mean taking the properties we are aware of in these cases to be actual properties of the objects external to us. The account of illusions I consider next involves making just this sort of claim, so I will wait to assess its merits.

external to us, his view amounts to denying the very presence of the phenomenal properties we are trying to explain. Needless to say, this seems like a non-starter.

Simply put, what Fish claims is not plausible. It is clear that in the McGurk case we experience [da], that in the binocular rivalry study with patchwork images we experience coherent images, that when the blind spot is filled in the visual experience includes that portion of the visual field, that in the flash and beeps case our visual experience is of two flashes, and so on. Even if we consider the case of a circular object being distorted by a lens so as to look elliptical, it is simply implausible to claim that there is no shape phenomenology here. We have as good a reason for thinking there is shape phenomenology here as thinking we are aware of the object's color or any other properties Fish would be happy to concede we are aware of in such an experience. Consider something like the McGurk effect, which we can all experience in the comforts of our own home. Here, Fish must claim that our processing of the incongruent inputs prevents us from hearing [ba] or lip-reading [ga], but also that we do not have the auditory experience of [da]. But this part of the experience is as salient to us as any of the things we accurately perceive at the same time. If he is going to deny this, we may as well deny that there are phenomenal properties at all.

But while I think it is clear that Fish has not offered a plausible explanation of mismatches, he may find my case question-begging. He may respond that I am simply appealing to introspection, which is misled by my false judgments and beliefs. Of course in the case of the McGurk effect it seems to me like I experience [da], but I am mistaken about my own experience. My mistaken belief does not lead me to experience something that is not there, but to merely take myself to be experiencing some such thing. So goes his retort.

Of course, by starting with the presumption that we can only experience phenomenal properties that correspond to actual properties of external objects we may think Fish is similarly guilty of begging the question. This starting point, though, leads to an implausible view of mismatches. I am not trying to defend the claim that we cannot be mistaken about the phenomenal properties that appear to us in experience. However, there are good reasons, aside from our obviously experiencing certain properties, to doubt Fish's explanation of mismatches.

In the snake case, it is plausible to think that we perceive the various properties of the rope and simply falsely judge that the object is a snake. It is at least plausible to say that we can have the same visual phenomenology without our mistaking the rope for a snake and that our actual mistake is not due to some additional snaky properties that appear. And even with the Kanizsa triangle, it might be plausible to say that we experience the figure and the properties of the incomplete circles—as Fish describes in the quote above—but no triangle or lines stretching between the shapes. Even here it seems plausible to say that we experience the phenomenal properties of the pacman shapes and of the white background, but simply misjudge that a triangle appears to us. But what about the Müller-Lyer case or the sorts of mismatches I have focused on? What are the phenomenal properties we can plausibly claim appear to us here? Here is Fish on the Müller-Lyer figure:

The perceptual experience of the Müller-Lyer lines has a number of phenomenal properties corresponding to various facts about the lines: their shapes, colors, relative locations, and so on. However, because of the way the lines have been contrived to produce misleading perspective cues, we cannot but passively deploy an inappropriate conceptual-recognition capacity—our capacity for recognizing things to be different in length—in our experience of these lines. Because of this, the subject

fails to become acquainted with the relational fact of the two lines' being the same length, and instead erroneously believes that the two lines are of different lengths. (p. 173-4)

What properties of the lines are we aware of yet we are not aware of their lengths? It is not clear how we can have an experience involving the lines' shape, locations, and "so on," yet not involving their lengths. What is such an experience like, what is its character? It would seem like if we are aware of the sorts of phenomenal properties Fish alludes to our experience will also include lines of a certain length, whether similar or dissimilar. Yet Fish denies we can be aware of their lengths as differing and also that we can be acquainted with their lengths as being the same.⁸¹

Things become even murkier when we turn to other cases. In the McGurk case, do we not actually have any auditory phenomenology? It is hard to see how we could be aware of some auditory properties yet neither have the experience of [ba], nor [da]. It is simply not clear what the character of such an experience would be, given that an auditory experience of [da] seems to simply involve the phenomenal properties corresponding to that syllable. In the binocular rivalry case with patchwork images, it is hard to see how there can be any properties we are aware of, given that our experience is of coherent images alternating, while it is not alternating images that are presented to us. In the flash and beeps case, do we correctly perceive the brightness of the flash but mistakenly judge that there were two rather than one? It does not seem correct to say that we can be acquainted with the actual properties of the flash, yet are not aware of there being a single flash or two flashes. And, regarding filling-in at the blind spot, Fish's account would lead him to say we simply are not acquainted with any phenomenal properties in the blind

⁸¹ Millar makes a similar argument concerning the interdependence of different features of our experiences. He relies on stronger claims about interdependence than I presume, though ones Fish seems to accept (see, for example, p. 146).

portion of the visual field and that our blind spot prevents us from being aware of the blue circle that is there. But the sorts of activity we see in V1 suggests that more is going on in the perceptual system than our simply not being aware of certain properties. And it is worth considering this perceptual system in a bit more detail.

For the sake of simplicity, let us think of our experiencing phenomenal properties as what results from the perceptual system completing its work. On Fish's view, when we are exposed to certain stimuli, our perceptual processing results in our experiencing certain phenomenal properties, but the processing also leads to our misjudging our experience and failing to become aware of certain properties of objects. On the alternative view I am defending, our perceptual processing actually results in our having the experience of [da] or our having the experience of red in the blind portion of the visual field. And when we look at the actual perceptual processing involved in various mismatches, we find that there is feedback to early areas of perceptual processing, and then further processing that takes place, finally resulting in the perceptual system completing its work. It therefore makes more sense to understand [da], or red in the blind spot, or the other properties we take ourselves to experience in cases where there are mismatches, as resulting from perceptual processing in just the same way that our experience of the pacman shapes in the Kanizsa figure result from our perceptual processing. Mind-independent objects have certain properties, but it is not as though we simply open the shutters and perceive them, or take in unambiguous input and straightforwardly process different features. Sensory input is noisy, input from different modalities interacts, and the very reason that certain stimuli can exploit certain processes is because those processes have to exist in the first place to facilitate our rapid perception of the external world based on the inputs we receive. The very same sort of processing that results in our experiencing the properties of the rope or the properties of the

incomplete circles in the Kanizsa figure also result in our experiencing the phenomenal properties we are seeking to explain in the cases of mismatches. On Fish's view, the perceptual system simply allows us to become aware of the properties of objects through causal interaction. Such a view might be plausible if the perceptual system operated in such a straightforward way, where various properties are unambiguously detected on the basis of neat input. It might then make sense to think that perceptual processing results in our being aware of various properties of the circles in the Kanizsa neon-colour spreading figure, but could not result in our being aware of a patch of neon blue that is not on the actual page. But the much more plausible view is that input is processed in a manner that results in our experiencing certain phenomenal properties that do not correspond to the actual properties of objects because—via the processing of the relevant input—our perceptual system takes these properties to be properties of external objects.⁸²

A second naïve realist strategy for trying to account for illusions is the converse of the one Fish pursues. According to naïve realism, perception involves our being directly aware of

⁸² We might think that Fish's view of illusions relates to recent discussions regarding the richness of consciousness. Recent experimental work has raised questions about just how richly detailed our perceptual experience is (see, e.g., discussions in Kouider et al., 2010; Cohen et al., 2016; Odegaard et al., ms). For example, though we may have the impression that our visual experience is detailed even in the periphery of our visual field, empirical findings cast doubt upon this: we do not notice when the letters outside our fixation point on a page are replaced by gibberish while we read (Grimes, 1996). These sorts of findings have prompted some to claim that we have an inflated sense of our perceptual abilities or that we can have a mistaken impression of our own perceptual experience. Somewhat similarly, Fish claims that in the case of illusions we are led to make false judgments about which phenomenal properties are actually present in experience. However, Fish's claim is not motivated by any empirical findings, and he does not support his view by pointing to things like the difference in visual acuity between the fovea and other parts of the retina—this last may be unsurprising, as he offers his account of illusions regardless of where in one's visual field the relevant objects are. Furthermore, the findings discussed in the context of debates about the richness of consciousness do not cast any doubt on our claims about what we experience when it comes to the mismatches I have discussed, whether because subjects are able to reliably report their experiences in such cases or because the phenomena concern foveal vision (as, for example, stressed by McLaughlin, 2016). Those raising questions about the richness of consciousness do not dispute that there are often mismatches between our experiences and the mind-independent objects that cause them.

mind-independent objects and their properties; thus, if objects do not have certain properties, we cannot be aware of such properties in experience. Hence, Fish's denial that certain phenomenal properties are present in the mismatch cases. The converse strategy would be for a naïve realist to accept that the phenomenal properties in question are present, but to argue that these actually are properties of the relevant mind-independent objects. This is the looks strategy.

Brewer (2008, 2011, 2017) offers such an account of illusions. Focusing on vision, he describes an illusion as “an experience in which a physical object, *o*, looks *F*, although *o* is not actually *F*” (2011, p. 102). While it is true that the object is not *F*, it is also true that the object looks *F*. His discussion focuses on three examples: a white piece of chalk in red lighting looks red even though it is white; a straight stick half-submerged in water looks bent even though it is straight; and, the lines in the Müller-Lyer diagram look to be unequal in length even though they are equal. The claim that the white piece of chalk looks red acknowledges the character of our experience when we perceive the piece of chalk—this experience includes the phenomenal property red. In each case, Brewer explains the object's looking *F* despite not being so by the object's having certain “visually relevant similarities” with objects that are *F*. Thus,

...a mind-independent physical object, *o*, looks *F* to a subject, *S*, in virtue of the fact that *S* is consciously visually acquainted with *o* from a point of view and in circumstances of perception relative to which *o* has visually relevant similarities with paradigm exemplars of *F*, where *visually relevant similarities* are similarities of the various kinds to which the physical processes enabling visual perception respond similarly, as a

result of both their evolutionary design and their development over the course of our lives. (p. 118)⁸³

On Brewer's view, an object has a certain look in some circumstance because in that circumstance it has certain similarities with other objects. Thus, in the Müller-Lyer case, the diagram, in normal conditions,

...has relevant similarities with a pair of lines, one longer and more distant than the plane of the diagram, one shorter and less distant; and those lines in themselves are a paradigm of inequality in length. (p. 102)

And, in the chalk case, a white piece of chalk illuminated with red light,

...has visually relevant similarities with a paradigm red piece of chalk, of just that size and shape. Their visually relevant similarity consists in the similarity of the light reflected from both. (p. 106)

The question is how exactly we are to understand these similarities. A rectangular sheet of paper and a rectangular piece of wood are similar in virtue of their actually having the same shape, and in perceiving each of them we will experience a rectangular shape. When it comes to illusions, things will understandably not be so straightforward. It is clear that the similarities will have to be in terms of things that are relevant for our experience. This means that similarities will have to consist in either properties of the objects themselves and how these properties interact with features of the particular circumstances, or in how certain stimuli are processed by our perceptual system. And this is just what Brewer claims:

⁸³ As we can see from this quote, Brewer understands perception as a three-place relation, where in addition to the subject and object we must consider various features pertaining to the context of perception, such as the subject's point of view and things like lighting conditions. The introduction of this third relatum is irrelevant for our present purposes, and does not make his account different from other versions of naïve realism in a way that is significant here. Brewer also understands perception in terms of acquaintance with objects, rather than their properties. Again, this difference is unimportant for our purposes, as we can understand similarity in terms of shared properties (p. 103). For ease of discussion I will put things in terms of properties.

Objects have visually relevant similarities when they share sufficiently many common properties amongst those that have a significant involvement in the various processes underlying vision. Thus, and very crudely, 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... (p. 103)

In the chalk example, a white piece of chalk and a red piece of chalk have the following relevant similarity: their surfaces are such that in different lighting conditions both can reflect similar patterns of light. Meanwhile, the horizontal lines in the Müller-Lyer diagram are similar to horizontal lines of unequal length in the following way: both stimuli are processed by our perceptual systems in ways that result in our having the experience of lines of unequal length. It is not the case that the processing of the horizontal lines in the Müller-Lyer diagram will be identical to the processing of unequal lines—this is obviously true for early stages of processing. Rather, the similarity between the two stimuli consists in their resulting in visual experiences that are similar. It thus seems like we can understand a look as a feature an object has in some circumstance due to some properties the object has. An object, o , that is not F will, in certain circumstances, look similar to objects that are F , because how o is will result in its reflecting light in a certain way or in its impacting our perceptual system in a certain way.

There are a number of variants of the looks approach to illusions, but they are all similar enough for our purposes.⁸⁴ Basically, they amount to claiming that in cases of illusion we do not

⁸⁴ For example, see Genone (2014). For discussions of various approaches, see Millar (2015) and French (ms). Martin (2010) offers an account of looks on which looks are mind-independent properties of objects. He does not apply his account to illusions, which is unsurprising given his approach to illusions. In general terms, his account is similar enough to the other variants, but

experience an object as having some property it lacks, but rather experience a property the object has, a particular look. A look is a property an object has in a particular situation. Thus, it is not that I experience a piece of chalk that is red; rather when the white piece of chalk is illuminated by red light it looks red. Its looking red is a property of the piece of chalk. My visual experience certainly includes the phenomenal property red, but I am not aware of the redness of the chalk, as the piece of chalk is not red. Its looking red in that situation is due to certain similarities it has with other objects.

Let us see how the looks approach can be applied to mismatches.⁸⁵ In the binocular rivalry study with patchwork images, the two images look like coherent images alternating. In the flash and beeps case, the single flash presented in that situation looks like two flashes. In the McGurk case, audio of [ba], when combined with lip movements for [ga], sounds like [da]. When we view the red square with a blue circle at its center from a certain angle and with one eye closed, it simply looks like a red square. Now, how do we explain these different looks—or, why is it that, for example, the single flash looks like two flashes even though it is only one? Why do we experience two flashes when there is only one?

One way of explaining how things look or sound (or feel or taste or smell) to us in mismatch cases would mirror how Brewer explains the chalk case. In the chalk case, the chalk looks red because in that situation it reflects light in the same way that a red piece of chalk in normal lighting would. It is this similarity that explains the chalk looking red. The character of our experience, the redness that appears to us, is due to how the piece of chalk reflects light,

given Martin's concerns it is particularly ill-equipped to handle illusions. Therefore, I will not examine it more closely.

⁸⁵ Once again, I am not sure whether Brewer and others who offer looks accounts will consider mismatches to be illusions, hallucinations, or some combination of the two. However, this is irrelevant to our present concerns. The question, here, is simply whether or not the looks strategy can help the naïve realist explain our experiences in mismatch cases.

which is due to the chalk's surface. It does not seem like such an account will work for mismatches, at least not most of them. Consider multimodal interactions. These do not occur because sensory signals targeting different modalities interact before reaching our sense organs. We experience two flashes or [da] because of interactions between the modalities that occur during perceptual processing. The auditory input in the McGurk case is not similar to the auditory input typically associated with our experience of [da]. And the visual input in the flash and beeps case is not similar to the input typically associated with our experience of two flashes.

Furthermore, in the case of the piece of chalk in different lighting, it is easy to understand how specific properties of the chalk—concerning its shape and surface—could result in different sensory signals reaching us depending on the lighting conditions. But the sorts of effects we see with mismatches are much more variable than that. It is not simply that things may look different in different lighting, but our experience of objects is sensitive to a variety of changes in the environment—concerning spatial and temporal relations to a variety of possible stimuli—as well as in our point of view and things like our eye movements. It is simply implausible to think we could explain the complex, diverse, and variable effects we see with mismatches in terms of properties of objects akin to the surface properties of the piece of chalk. This would require attributing countless properties to objects in order to explain the possible ways they may look or sound in different situations.

A second way of explaining how things look or sound (or feel or taste or smell) to us in mismatch cases would mirror how Brewer explains the Müller-Lyer illusion. In this case, the horizontal lines look to be unequal in length because of how they impact our perceptual system in that situation, i.e. flanked by the different hashes. The character of our experience—the lengths the lines appear to have—is due to how our visual system responds to the stimuli. The

way our visual system responds, the story goes, is similar to how it typically responds to lines of unequal length. Once again, this is supposed to be a property of the horizontal lines in the Müller-Lyer diagram. In the McGurk case, this means that the audio of [ba] presented in that situation causes us to have the experience of [da], and that this is to be understood as a property of [ba]. The stimulus has the (auditory equivalent of a) look it does in that situation because the way our perceptual system processes it results in our having an experience that is like the experience we normally have when our perceptual system processes audio of [da]. But there are at least two problems with such an explanation. First, and most importantly, it seems to mistake what is actually doing the explanatory work. Second, it is not at all the case that the perceptual system handles the different stimuli in the same way, at least not in a sense that helps the looks theorist's case. I will start with the second point.

In the flash and beeps case we experience two flashes. According to the looks approach, the flash—when experienced in that situation—looks like two flashes.⁸⁶ The way this flash is similar to two flashes has to do with how the flash is processed by our perceptual system in this situation. Specifically, the visually relevant similarities between the flash and two flashes will concern this perceptual processing. As quoted above, Brewer notes that these similarities “are identities in...the way in which stimuli are handled by the visual system.” But while the resulting experiences may be identical, it is simply not true that the stimuli are handled by the visual system in identical ways. For starters, the initial sensory input differs, which necessarily means

⁸⁶ It may make more sense to consider the setting more generally, the stimuli present, or the scene, rather than focus on individual objects. Doing so may create problems for Brewer's specific form of naïve realism, but even if so it would seem to not be an issue for other naïve realist views. It simply strikes me as more natural to understand many of the experiences in mismatch cases as resulting from our perceiving a scene or arrangement, rather than as perceiving a specific object or even groups of objects. However, my argument here does not hinge on this issue at all. Also, note that while I discuss the problems with this explanation in the context of mismatches, the same points apply to the Müller-Lyer case.

that the processing taking us from the input to our experience will be different. There is a sense in which the stages of processing will be similar, or there will be similar cross-modal integration, or similar feedback processing, and so on. But in these senses it is trivially true that all stimuli are processed in similar ways. It is not until certain processing takes place in the flash and beeps case that does not take place in the typical case that the processing in the two cases will continue in a similar, or possibly even identical, fashion. The same is true for all of the other mismatches discussed, whether various forms of visual completion, or temporal filling-in, or geometrical misperception. The visual processing of the one flash in the flash and beeps case and of two flashes in a normal case may be similar only in the latter stages or visual processing. But this is the sort of datum one would point to in order to support the explanation offered in (iii), which invokes mental representations and the brain. And this brings us to the first problem with this way of explaining looks.

According to this approach, we explain our experiences by pointing to a property of the object, but that property is just to cause a certain response in us. The red square with a blue circle at its center is processed in such a way in a particular situation—when viewed from a certain distance and angle with one eye closed—that it results in our experiencing a red square. Why does it look this way? Because, it is alleged, the stimulus is handled by our perceptual system in a similar way as a red square is handled. But our perceptual processing of the two stimuli is similar only in latter stages of processing. In particular, the responses from our perceptual system to the two stimuli will be similar when we look at the result of the processing. Given this, it is hard to see why we should consider it a property of the mind-independent object that it causes such a response. Doing so is to look in the wrong place for an explanation of our experience. The reason the two stimuli cause similar experiences is because their processing results in similar

brain states, which subserve the experiences. The brain states certainly result from our causally interacting with the stimuli, but the reason why our perceptual system responds in this way, and the reason why the experiences have the same (or similar) character is to be found by looking at the perceptual system. It is not clear what sort of properties the looks theorist is claiming mind-independent objects have, or if it makes sense to think of these as genuine properties of objects. The individuation of properties in this case is made on the basis of the perceptual system and its responses. Thus, to claim that the explanation for the character of our experiences is to be found in the mind-independent object is ad hoc, is to offer a just-so story.

The issue here is how we should explain the character of our experiences in mismatch cases. Is the phenomenal character determined by a representation generated by the brain, or by mind-independent objects and their properties? The looks theorist tries to argue that we should understand the phenomenal character of our experiences as being determined by mind-independent objects. Brewer's view "is that visually relevant similarities are those that ground and explain the ways that the particular physical objects that we are acquainted with in perception look" (p. 103). How is it that the two patchwork images are similar to the two coherent images? Well, they are similar in that, in the appropriate situations, both can result in the same experience. The question we want answered is why this is the case. By looking to perceptual processing and the brain, we find a good explanation: the brain generates a representation, which determines the character of our experience. Looking to the mind-independent object does not provide a plausible explanation. The claim that the object has the property, in that situation, of looking a certain way because it has the property of impacting our visual system in a certain way is not to explain how or why it is that the patchwork images presented in binocular rivalry result in our experiencing the coherent images alternating. It is

simply to tell us what we already know: that the stimuli have that effect. The explanation for why that is certainly includes details about the objects and events external to the individual, but is fundamentally a matter of what is going on inside the individual. Consider filling-in of the blind spot. It is not the stimulus but the visual system that fills in that part of the visual field, resulting in our having the experience we do. What is filled in—so what we experience in that part of the visual field—depends in part on the properties of the stimulus, but not some special properties or the property to cause some response in us. It is properties of the surround that play a role. The phenomenal properties we experience in the blind spot can be traced to the brain's active filling-in. The same is true for other forms of visual completion and for temporal filling-in, and the brain's role is similar in the other mismatches.

Since the naïve realist accounts of illusions do not provide a plausible explanation of our experience in mismatch cases, we are right to accept (iii). We are justified in invoking mental representations, as doing so offers the best explanation of our experience in these cases.

But before moving on, we may think there is another way one could respond to (iii) that merits consideration. French and Walters (forthcoming) argue that the traditional argument from illusion is invalid. As we saw in chapter 1, the argument moves from the claim that in experience we are aware of something that is *F* to the claim that we are not aware of the mind-independent object that is not-*F*. Sometimes this is put, instead, in terms of what we are directly aware of: thus, we are not directly aware of the mind-independent object, but instead are directly aware of something else (e.g. a sense-datum). The argument then attempts to generalize this conclusion to all cases of perception. However, French and Walters argue that the claim that we are not directly aware of the mind-independent object in the case of illusion does not follow from the preceding premises and is therefore invalid. We are entitled to claim that we are directly aware

of something that is *F*. But, they say, this does not rule out that we are also directly aware of the mind-independent object that is not-*F*.

We may think their point relates to the mismatch argument and the argument that mismatches are ubiquitous in the following way. Consider (iii):

(iii) We can explain what the individual experiences if we suppose that the brain constructs a representation that determines the character of their experience. The individual is aware of this representation. We cannot explain the character of the individual's experience simply by pointing to the external objects, since their experience, or what appears to them, differs in the way it does from the external objects. In this case, the individual is not immediately aware of the external objects.

Given their point about the argument from illusion, French and Walters may take issue with the last sentence of (iii). While not disputing that in these cases we are immediately aware of a representation generated by the brain, they may claim that this does not rule out that we are also immediately aware of the external objects. But whatever the merits of their argument regarding the validity of the argument from illusion, such a claim should not be taken seriously here.

For one, given what "immediately aware of" means, it is not even clear how we should understand the claim that an individual is both immediately aware of a representation generated by the brain and the objects external to them. But, more importantly, there is absolutely no reason to claim that an individual is immediately aware of the external objects in these cases. We explain the character of the individual's experience with the representation. The external objects serve no explanatory role here. To insist that we are nevertheless directly aware of external objects here would be akin to claiming—to borrow an example from my first philosophy professor—that the reason a door opens and closes smoothly has to do with the various physical

forces at work, but there are also invisible and inert angels there to supervise the process. Of course, given a certain understanding of angels, it may be impossible to disprove this theory of mechanics. But the only reason we would have for supposing angels were present would be our commitment to some theory claiming their existence. We can explain what we want to explain—the movements of the door—entirely in terms of various laws governing matter.⁸⁷

5. Perceptual Processing and the Brain

The next step in the mismatch argument is:

(iv) Perceptual experience depends on brain mechanisms that always function in the same way.

If the character of our experience in certain instances is determined by a representation constructed by the brain, then the character of our experience is always determined by a representation constructed by the brain.

This premise generalizes the conclusion about mismatch cases to all instances of perceptual experience. In the mismatch cases, a representation is generated by the brain, and this representation determines the character of our experience, determines the sensory qualities present in experience. Here, the claim is that the brain mechanisms involved in perception always function in the same way: such a representation is always generated and is involved in

⁸⁷ Hickerson (2004) may offer a response that is similar to the French and Walters claim. He argues that Smythies and Ramachandran's claim that naïve realism is refuted by the binocular rivalry study with patchwork images is misguided. He claims that Smythies and Ramachandran's target can simply deny that images constructed by the brain are phenomenal objects, and instead assert that the only phenomenal objects are the images C and D. As regards the mismatch argument, this would be to claim that we are directly aware of the two patchwork images. It is not clear why Hickerson thinks the naïve realist is able to make such a claim. If the naïve realist is to maintain that in the cases of mismatches we are directly aware of external objects, they must explain how this is so. Hickerson argues that such an explanation is possible, but does not suggest what it may look like. We can understand the discussion in this section as demonstrating that different explanations that the naïve realist may try to offer are unsuccessful. It may be that Hickerson has in mind something like the French and Walters claim.

this way. In this section I defend (iv). I argue that we are right to think that perceptual processing always operates in this same way. My focus in this section is solely on the mismatch argument, so we can set aside the issue of how common mismatches are in our day-to-day lives. Here, we need focus only on a handful of mismatches, considering only things like the McGurk effect, the flash and beeps case, or illusory contours—mismatches elicited in the lab or with specially constructed stimuli.

The idea behind (iv) is that perceptual processing in the brain always functions in fundamentally the same way. The brain receives input, which initiates a series of stages of processing, culminating in an internal representation of the scene, which determines what we experience. The stages of processing are at a general level always the same, involving the same sorts of interactions between brain areas and feedback and feedforward pathways. This is not to say that different sensory inputs cannot result in certain processes or pathways being more active in certain instances. For example, multimodal integration, which takes place all the time, may result in visual information strongly influencing audition in certain cases but not others. But these differences explain how different inputs may result in different outputs, where outputs are understood to be the result of this processing, i.e. what we experience. Later stages of perceptual processing represent the external scene, as discussed in the previous section. The precise neural activity here, corresponding to particular information or properties, is due to the sensory inputs and the precise way in which these were processed by the brain. But in all cases this activity is a representation of the external world, which is related to our perceptual experience in the same way. This view of perceptual processing is not due to a single finding or study, but is supported by our general understanding of the brain and its role in our perception of the world, as well as

how it subserves our perceptual experience—an understanding resulting from extensive scientific investigation.

Of course, the part of this processing that we are most concerned with here is that the brain generates, or constructs, a representation that determines the character of our experience. Since this happens in mismatch cases, we should think that this happens all of the time; it is implausible to suggest that the brain generates a representation in some cases but not in others. As noted when I first presented the mismatch argument in the previous chapter, and as Smythies and Ramachandran note, the brain receives sensory inputs—it does not have direct access to what causes those inputs. In setting up the binocular rivalry study with patchwork images or the incongruent stimuli in the McGurk case, the experimenter or observer knows what is being presented to the individual.⁸⁸ But the individual presented with the stimuli has the experience they do because of how the brain processes its sensory input in these cases. In the McGurk case, the brain receives the auditory and visual input and begins to solve the problem of determining what external objects and events caused that input. Multimodal integration happens in the course of normal perceptual processing, and here, because the information in these different modalities is incongruent, this information influences processing in a certain way. We can understand this as the brain interpreting what is external to us, or making sense of the input in a certain way.⁸⁹

The same thing happens when input is not drastically incongruent in this manner: the brain

⁸⁸ We know what stimuli the individual is presented with on the basis of other experiences. For example, the experimenters may view the images in non-rivalry conditions or experience the auditory and visual stimuli in isolation. There is no assumed directness when it comes to the experiences of the experimenters or other observers, as alleged by Wilcox and Katz (1984). Although primarily concerned with slightly different issues than the ones I focus on, Wilcox and Katz argue that an assumption of directness undermines attempts to bring certain empirical work to bear on philosophical questions relating to perception. For a reply to Wilcox and Katz that focuses on their main concerns, see Wright (1986).

⁸⁹ As for what sort of mechanisms or computations may underlie such a process, that is a separate issue, one I will touch on briefly in the next chapter.

receives input and tries to solve the problem of determining what external objects and events caused that input. The result of this processing is a representation of those external objects and events. In the McGurk case, our experience is of [da] because the brain's computations determine that to have been the stimuli that resulted in the sensory input received. The same happens when I have the experience of my nephew saying "dada."

One of the reasons it makes sense to think that perceptual processing always operates in the same way is the fact that sensory input is always noisy. This may not be apparent from our perceptual experience, but it is important to keep in mind. The fact that input is noisy in normal circumstances means that, in important respects, the everyday case is not different from the mismatch case. In all cases, the starting point for the brain is noisy and often ambiguous input. Speaking very roughly, we can have input caused by v that seems like noisy input caused by y , and we can have noisy input caused by z that actually seems like less noisy input caused by w . The first is the McGurk case, and the second would result in a mismatch and our experiencing w . Furthermore, as discussed above, while experimental conditions are highly controlled and may strike us as extremely artificial, they often mimic the sorts of stimuli we confront in our day-to-day lives.

Like the spreading premise in the traditional arguments from illusion and hallucination, (iv) generalizes a conclusion about certain cases to all cases of experience. However, as I noted in the previous chapter, there are important differences between (iv) and the spreading premise.

The spreading premise states that the same account must apply to both veridical and illusory or hallucinatory experiences. The reason why the same account must apply is that these different experiences are subjectively indistinguishable. From the first person perspective an illusory experience need not stand out in any way compared to a veridical experience. Given that

they are subjectively indistinguishable, it would seem that they have the same fundamental natures—for, if they had different natures, one would think this would lead to their being different in ways apparent to introspection. A second reason offered for why the same account should be given is that illusory, hallucinatory, and veridical experiences can form a continuous stream. But this is closely related to the subjectively indistinguishable point. If the different cases were subjectively distinguishable, then the stream of experience would be discontinuous, and it would be evident when there was a shift from one sort of case to another. That my experience while undergoing various illusions segues fluently into non-illusory experiences is a testament to their common natures.

The reason for generalizing to all cases in the mismatch argument is different. We generalize because the same thing is going on in all cases of experience, the same neural processes are involved. This gives us good reason to think that a representation is constructed in all cases. This is not to say that mismatch cases are not subjectively indistinguishable from or that they do not form a continuous stream with non-mismatch cases. Nor am I claiming that such facts do not suggest that they have the same fundamental nature. But the reasoning behind (iv) is stronger than what supports the spreading premise. According to (iv), we generalize because, when it comes to the processes that underlie our perceptual experience, the same thing is happening in all cases. This does more than merely suggest that these cases all have the same nature. Rather, this tells us why we should understand and explain them all in the same way. This difference proves important in the face of replies.

The naïve realist response to the spreading premise is disjunctivism. The naïve realist argues that despite the experiences being subjectively indistinguishable, they are not of the same fundamental kind, and therefore should not be given the same account. Whether or not we find

this manoeuvre convincing, there is certainly space there for the naïve realist to make it. They can agree that the different experiences are subjectively indistinguishable, but disagree about what follows from this.

Given the similar roles of the spreading premise and (iv), it is reasonable to think the naïve realist would try the same manoeuvre in response to (iv). But such a move does not seem to be available here. The reasoning behind (iv) seems to undercut any force such a response may have. If the naïve realist accepts that the same processes are involved in all cases, that perceptual experience depends on brain mechanisms that always function in the same way, then it does not seem to matter if they deny that these cases are all of the same fundamental kind. Whatever it would mean to be of a fundamental kind, here, would be divorced from the question of how we explain an experience. Basically, (iv) gives us a good reason why we should explain all experiences in the same way. Because it does this, the disjunctive move does not have any traction.

Instead of conceding that perceptual processing always operates in the same way, one may try to deny this claim. This would be to claim that while a representation is generated by the brain in mismatch cases, such a representation is not generated in non-mismatch cases.

The problem with this response is that we have absolutely no reason to think that it is correct. We have no reason to think that perceptual processing operates a certain way in mismatch cases and a different way in other cases. There is no reason to think the brain is haphazard in this way. But, more importantly, the brain does not know what is causing the sensory input. To return to the rivalry study: my brain does not know whether I am presented with images A and B or with images C and D. So the claim that my brain generates a

representation only in the latter case is untenable. The same can be said for all other cases when mismatches occur.⁹⁰

Finally, one may reply that while a representation is always constructed by the brain, such a representation does not always determine the character of our experience. Specifically, the response goes, in cases of veridical perception it is mind-independent objects that determine the character of our experience.

But, once again, we have absolutely no reason to think that this is correct. Just as perceptual processing operates in the same way in all cases, the relationship between the representation generated by the brain and our perceptual experience is the same. It is hard to understand how it could be that such a representation determines the character of our experience in some cases but not in others. As just noted, the brain does not know whether the processing of some input will result in a mismatch or not. And the actual presence of the requisite mind-independent objects and properties cannot be the crucial factor here—to think so would be to have things backwards. The idea that a representation “does not need to” determine the character of an experience in some case because the external objects can not only assumes that naïve realism is correct, but requires that the decision—if you will—about what determines the character of an experience is made only after the character of an experience is compared with the mind-independent objects present, in order to determine if there is a mismatch. Like the French and Walters proposal regarding what we are immediately aware of in the case of illusions, the

⁹⁰ For a similar point in the context of the causal argument for the common kind assumption, see Robinson (2013) on the “non-arbitrariness” of the hypothetical hallucinations. In the context of that argument, which I discuss in the first chapter, Robinson takes the non-arbitrariness of the brain’s functioning to support the principle of “same proximate cause, same immediate effect.” Above, I explain why we should think the brain always generates a representation. The reasoning at this stage of the mismatch argument may remind us of the causal argument in that a certain regularity in the brain’s functioning is appealed to. The next point I make is in the same vein.

mind-independent objects do not serve any explanatory role in non-mismatch cases. We can already explain the character of our experiences in such cases by pointing to the representations generated by the brain. To insist that we are instead, or also, immediately aware of mind-independent objects and properties is baseless.

Thus, we accept (iv). We are right to think that perceptual experience depends on brain mechanisms that always function in the same way; namely, that the brain constructs a representation that determines the character of our experience.

In the previous section, while considering mismatches, we may have anticipated the following sort of objection: in such cases the mismatched portion of our experience has its character determined by a representation constructed by the brain, but the character of the veridical portions is determined by the external objects present. Brown considers such a reply when articulating his argument, discussed in the previous chapter.

One might maintain that the perceptual display is a heterogeneous compound of elements of the external arrangement and elements generated by our perceptual system. But given the coherent integration of the various elements of the perceptual display, we get a more intelligible picture if we consider the entire perceptual display to be a brain-construct that is numerically distinct from the external arrangement involved in its causation. (p. 48-9)

I agree with Brown that, quite simply, the more reasonable view is that the character of the whole experience is due to the brain. However, the reasoning in this section helps us further support this position. All of the points I have made about different experiences can be made about different aspects or portions of a single experience. Therefore, the objection need not trouble us.

6. Most Cases are Good Cases

The argument that mismatches are ubiquitous is inspired by the empirical work examined earlier in this chapter. This work reveals that mismatches are ubiquitous in our everyday lives. In defending (iii), I argued that when it comes to mismatches the character of our experience is determined by a representation constructed by the brain. That mismatches are ubiquitous means that representations being involved in experience in this way is the norm. This is (iv*):

(iv*) Such mismatches between the objects and events external to the individual and the individual's experience are the norm. If the character of our experience in these instances is determined by a representation constructed by the brain, then the character of our experience is normally determined by a representation constructed by the brain.

This would seem to be a fatal problem for naïve realism, which denies that perceptual experience involves representations in this way. If it is the norm that the character of experience is determined by a representation constructed by the brain—and therefore not by mind-independent objects and their properties—then it seems we must reject naïve realism.

Whether or not mismatches are the norm is a relatively straightforward empirical question. And the empirical work discussed above, in section 2, supports an affirmative answer. By examining the different situations in which, and the different ways in which, mismatches occur, this work shows that mismatches are the norm in day-to-day life.⁹¹ Adding together the mismatches resulting from the different sorts of processes discussed, the proportion of our experiences involving at least some minor mismatch—or illusory—element or aspect is

⁹¹ Moreover, it seems reasonable to think that further empirical work—in particular in the non-visual modalities—will provide even more evidence of mismatches that can occur in our daily lives and shed light on various ways that the relationship between our perceptual experience and the mind-independent objects and events causing that experience may vary and come apart.

significant. So the naïve realist cannot simply reject (iv*). In response, if they hope to rescue their theory, they must say the following: *What the empirical work shows is that there are many more mismatches than we previously thought. That is certainly interesting, and it means that many more experiences will be classified as bad cases than we previously thought. But we can preserve our analysis for the good cases and maintain that perception involves the direct awareness of mind-independent objects and their properties.* However, there are two independent reasons why this is not a good response.

A philosophical theory of the nature of perception strives to explain our experience. A theory casting some experiences aside as bad cases, and as not amenable to the analysis offered by that theory, may not be such a problem. Though we may prefer that a theory be able to account for all of our experiences, if bad cases are rare—or even merely possible but non-existent—then we may think it fine to exempt them from our analysis. But if bad cases are the norm, and if the analysis of perception offered by the theory does not apply to the bad cases, then we have a theory that only applies to a small number of cases. And, at this point, it is not clear what good this theory is. It is one thing to learn that things are not quite as they seemed, that perceptual experience involves many more mismatches than previously thought. It is quite another to learn that a theory of perception applies only to a shrinking number of experiences. In the latter case, we are left with a theory that purports to explain only a small proportion of the phenomena of interest. Not only is such a theory of little use, but it appears wrongheaded.⁹²

⁹² The issue here is not that these experiences are labeled as illusions, or hallucinations, or mismatches, or whatever else. The issue is that the analysis offered by the theory does not apply to these experiences. This means that the naïve realist does not have a theory of experience. Furthermore, since illusions are typically thought of as perception—or at least misperception—the naïve realist does not even have a theory of perception. And this is bad. The theory simply does not cover what we want a theory to cover. Others make a similar point in slightly different contexts. Robinson (1994, p. 159), Smith (2002, p. 28), and Millar (2015, p. 611) argue that if

This problem is particularly acute for naïve realism, given that the theory is alleged to capture our common sense view of perception or, at the very least, how things seem to us when we introspect on our experience. Let us grant, for the moment, that the naïve realist is right about how things seem to us in experience: it seems like mind-independent objects and their properties are directly present to us in experience. The primary motivation for naïve realism is that it does justice to this. But it does not merely sometimes seem to us that mind-independent objects are directly present in experience. Barring drug-induced hallucinations, psychotic episodes, and other rare occurrences, the phenomenology of our experience is fairly consistent. In mismatch cases, it seems to us like mind-independent objects and their properties are directly present in experience. But things are not actually as they seem—and this is far from a rare occurrence. So why should we put any stock in how things seem to us when we introspect on our experience when it comes to deciding which view about the nature of perception is correct? The naïve realist promises to offer an account of experience that meshes with how things seem to us commonsensically or when we introspect on our experience; but by casting aside mismatches as bad cases the naïve realist's account no longer meshes with how things seem much of the time. Since the theory's agreeing with how things seem was the motivation for adopting the theory, and since the theory no longer agrees with how things seem, we no longer have any reason to hold on to the theory, even if only to account for the good cases. Learning that mismatches are so common and seeing how we explain our experience in these cases undermines the consideration that was motivating the acceptance of naïve realism.

the naïve realist analysis is not meant for illusions and illusions are pervasive, then we should reject the theory. But the phenomena discussed in these contexts are not things I have argued are mismatches. For example, Millar points to inaccuracies in our perception of how far objects are from us. Whether we find these arguments compelling will depend on whether we consider the phenomena to be illusions and if the case has been made that they are pervasive.

The second problem with the naïve realist's response to (iv*) is that most cases are perfectly good cases of perception. Pre-theoretically, our perception of the world seems to be pretty good. For the most part we get on in the world pretty well and our experiences form a coherent sequence. It is true that mismatches are the norm. And there will certainly be cases where the discrepancy between our experience and the external world is drastic in the ways we sometimes see in laboratory settings. But for the most part in our day-to-day lives mismatches will involve only a small part of some experience, a single element, portion, or aspect of an experience. And within the flow of our experience of the world these will often be things we do not notice, or attend to, or even form beliefs on the basis of. Consider, for example, how vertical lines appear longer to us than horizontal lines of the same length or how multimodal interaction can shift where a sound appears to come from. Now consider walking down a city block in search of a particular address and how an experience containing such a mismatch is immediately replaced by another experience and another. Arriving at one's location, our sense of the walk is similar to one's sense of a story we have read—we retain the gist, but not every detail, especially not the mundane ones. We get to where we were trying to go with little apparent effort. What McLaughlin says about vision goes for perception in general—once again:

The normal visual geometric misperceptions that I've highlighted don't pose a practical problem. There's much that vision gets right. Needless to say, relying on vision, we do just fine behaviourally, and we do spectacularly well cognitively. (p. 291)

Not only do we get on in the world just fine, but we do so because of how our perceptual system functions. It is the normal functioning of this system that allows us to get on in the world and gives rise to mismatches. We may be right to consider some of these mismatches illusions, or even hallucinations, but in most cases the mismatch will be minor and of no consequence. In

these cases, we should hold on to the idea that most experiences are perfectly good instances of perception.

The lesson is not that perception is wildly inaccurate and that it involves mostly bad cases, but that we need to rethink what it means to be a good case. In philosophical discussions of perception the contrast is made between veridical and non-veridical experiences.⁹³

Hypothetical hallucinations may be considered, as well as well-worn examples of illusion, like a white wall in red light, or a stick half-submerged in water, or a coin seen at an angle or through a prism, or even the Müller-Lyer illusion. But often little is said about how exactly we are to characterize veridical perception. If illusions involve getting it wrong, or experiencing some object as having a property it lacks, the implication is that veridical perception involves getting it right, or experiencing objects as having only properties they actually possess. But there may be difficulties with the idea of correspondence that seems to be suggested here (Schwartz, 2016). And the empirical work seems to put pressure on such a notion of veridical perception. If “veridical” is used in contrast to illusions and hallucinations, which are seen as bad cases, then it seems we need to take a closer look at this notion in light of the empirical work relating to mismatches. I will explore this issue in the next chapter.⁹⁴ For our present purposes, the key point is that we should not consider all mismatches to be bad cases.

⁹³ There may be further sub-divisions, and also different ways of classifying veridical cases, illusions, and hallucinations—e.g. whether illusions are to be considered cases of (mis)perception.

⁹⁴ I am not claiming that experiences where there are mismatches are completely accurate. Rather, if mismatches are the norm and these are often good cases, we may want to (re)examine our notion of veridical perception. Note that whether we consider some mismatches as cases of veridical perception does not mean that the naïve realist analysis can suddenly be applied here. The naïve realist cannot adequately explain mismatches. This remains true regardless of how we classify them.

Therefore, mismatches being the norm means we should reject naïve realism. The naïve realist's response to (iv*) does not save the theory. We have two reasons to dismiss the response; on its own, each is sufficient to reject the reply. Based on recent empirical work, we conclude that the character of our experience is normally determined by a representation constructed by the brain. Naïve realism is refuted. And this conclusion supports the sort of indirect realism articulated in the first chapter.

Now, we may wonder whether we are entitled to move to the conclusion of the argument that mismatches are ubiquitous, or if more needs to be said. That conclusion is:

(v) Therefore, in perceptual experience we are aware of a mental representation generated by the brain; we are never aware of mind-independent objects and their properties.

We may think that (iii) combined with the ubiquity of mismatches is enough to refute naïve realism and support indirect realism, but that, nevertheless, moving from “normally” to “always” is a bit quick.

The ubiquity of mismatches is significant here. We are not trying to move from a handful of odd cases to all cases. Rather, we have very many everyday cases where perception involves mental representation. It is reasonable, then, to think that in perceptual experience we are always immediately aware of a representation constructed by the brain, and never mind-independent objects and their properties. We have no motive to deny this, no reason to think that mismatch cases and non-mismatch cases differ in their natures. But even if mismatches are the norm, and even if the gap between this and all cases is much smaller than between some and all, we may want something more conclusive. If so, we can add (iv) from the mismatch argument, which I defended in the previous section. There may be other equally compelling considerations that can

help us bridge this gap—a much smaller gap than the one in the mismatch argument—but since I already defended (iv) we can deploy it here.

7. Conclusion

In this chapter I argued, based on recent empirical work, that we should reject naïve realism.

I defended the mismatch argument, which claims that in perceptual experience we are aware of a representation constructed by the brain, a representation that determines the character of our experience. We are not, as the naïve realist claims, directly aware of mind-independent objects.

I also defended the argument that mismatches are ubiquitous, which establishes the same conclusion.

In the first chapter I argued that whether or not perception involves mental representation is a particularly important question, one that constitutes the disagreement between naïve realism and indirect realism. Naïve realism denies that perception is a matter of mentally representing the world, such that the character of our experience is determined by a mental state and its properties. Indirect realism, in contrast, takes perception to involve mental representation in this way. Together with critical realism and intentionalism, it forms a family of views that stand on the same side of the representation question. In the next chapter I will consider whether the empirical work helps shed any light on the disagreements within this family of views. But, here, progress has been made. Recent empirical work on perception helps us answer our question: perception involves mentally representing the world, such that the phenomenal character of our

perceptual experience is determined by a representation. Thus, naïve realism is false, and this work supports indirect realism.

Chapter 4

Concluding Remarks: Veridicality, Prediction, and Implications

1. Introduction

I have argued, based on recent empirical work, that perception involves mental representation and that the character of our perceptual experience is determined by an inner representation. In this closing chapter I briefly explore a few different issues related to this conclusion and the arguments offered in support of it. First, I will examine the notion of veridical perception in light of the fact that mismatches are the norm. I will then look at whether the conclusion I defend is implied by theories that take the brain to be an organ for prediction error minimization, as argued by Hohwy (2013, 2016). Finally, I will touch on the disagreements between different views that take perception to involve mental representation and on some of the philosophical questions that remain, despite the progress made.

2. Veridical Perception

Recent empirical findings concerning a range of perceptual phenomena indicate that mismatches are ubiquitous in everyday life. Some mismatches may very well be considerable, but in most cases there will be only subtle discrepancies between how things appear to us in experience and the external objects and events causing those experiences. Given how common mismatches are, lots of normal perception we are happy to consider perfectly good turns out to include minor mismatches. We may find this surprising. While we have all experienced illusions at one point or another, the vast majority of our perceptual experience seems to provide an accurate guide to the world we inhabit.

Discovering that lots of experiences that seem like perfectly good cases of perception actually involve some mismatch element or aspect may be jarring. Faced with the ubiquity of mismatches we seem to have two options regarding the veridicality of our perceptual experience. We could say that much less experience than previously thought is actually veridical. Or, we can adopt a looser notion of veridical perception and retain the idea that most perception is good, veridical perception. I will argue that we should do the second.⁹⁵

2.1. Veridicality, Illusions, and Constancies

Schwartz (2016) examines the notion of perceptual veridicality, in particular as it relates to visual constancies and illusions. He notes how the standards of correctness for perception can vary and argues for a pragmatic approach to veridicality. His discussion will help to guide us as we consider mismatches.

As Schwartz notes, constancies are typically taken to be cases of veridical perception. When we experience an object's properties as constant despite changes in viewing conditions and retinal stimulation, we are thought to perceive the object as it really is. Meanwhile, failures of constancy are often taken to be cases of illusion. We are said to suffer an illusion when viewing conditions lead us to experience an object as having a colour or shape it does not actually have. The notions of constancy and illusion, therefore, both depend on that of veridical perception.

Perceptions are typically said to be veridical, when they *appear phenomenally* to be in agreement with the physical environment.

⁹⁵ Note that what I argue here has no bearing on my arguments against naïve realism in the previous chapter. Naïve realism cannot account for mismatches; this remains the case whether mismatches are classified as veridical cases, illusions, or in some other way.

Veridical experiences accurately reflect/mirror how things are in the world. (p. 384)

Illusions are taken to be cases where our experience deviates from how things are in the world. Constancies are taken to be cases where, despite changes in viewing conditions, our experience agrees with, or corresponds to, how things actually are. But because of complexities relating to the veridical/non-veridical distinction, Schwartz argues, ambiguities arise regarding constancy and illusion.

His discussion focuses on basic spatial properties—size and shape—and colour. I will ignore colour, so as to avoid the question of objective reality as it relates to colours, and focus on size and shape.

Consider size perception of the 250-foot-tall building... If veridicality lies in experiencing the building as having this metric size or even nearly so, our perceptual experiences of its size are likely to be nonveridical. If all it takes to be correct is that the building is presented as more than 35 feet tall, as bigger than I am, or in any number of other imprecise ways, all such experiences will be veridical. Demanding precision renders one assessment of accuracy, adopting looser veridicality standards another. In addition, there are an unlimited number of more precise and less precise characterizations of the building's size that visual experience can or can fail to correspond.

Choosing among alternative standards of correctness, though, is only half the problem in assessing perceptual veridicality. The other half is determining what the experience *purports* to represent for the perceiving subject. (p. 384-5)

These same issues exist when it comes to evaluating size constancy. My perception of the building's size may not change from when I am standing 30 feet from it to when I stand 60 feet from it. But this does not mean that my assessment of its actual size will be accurate in either case, given certain standards of correctness. We find the same ambiguities with comparative size perception in everyday life. A parking space may look to be more or less the same size from different vantage points. In all cases, say, it will look larger than 15 feet long and shorter than 20. But we may incorrectly think it is too small for our car until we align our car with it. Similarly, we only realize that the step stool is not tall enough once we have moved it under the light fixture. But that does not mean our experience of its size has changed or that one of our experiences is veridical and the other not.⁹⁶ Our perception of size, including comparative size, is often inaccurate, yet we are not inclined to call all such experiences illusions. And our experiences can be inaccurate even if constancy is exhibited.

Even though assessments of shape are typically more accurate than of size, similar issues exist regarding shape perception. For example, two tables of identical size and shape can be placed at different angles such that they appear to have quite different shapes. As Schwartz notes, it is not clear which table, if either, is seen accurately here. Moreover, if seen individually, we are not inclined to consider our experience of either table inaccurate; and, our perception of the table may exhibit constancy when viewed from different perspectives. Meanwhile, if presented with three shapes in an array, two of which are identical save for orientation, we are

⁹⁶ These examples of inaccurate size perception are Schwartz's—for discussion of further examples, see p. 384-9. Note that my objective here is simply to present Schwartz's discussion and the issues he raises. I am not claiming that these cases are mismatches or illusions.

often unable to judge which two are the same. Yet, we are unlikely to think our experience of any of the shapes is non-veridical, in particular if viewed separately.⁹⁷

By drawing our attention to the ways our perception of objects may often be inaccurate, Schwartz is not trying to argue that our experience is rarely veridical. Rather, his point is that it is often unclear whether or not an experience is veridical because the notion of veridicality itself is unclear. He thinks that,

...at root, many of the puzzles concerning perceptual veridicality and the notions defined in terms of it are traceable to the murky idea of “appearances agreeing or disagreeing with reality.” (p. 399)

When it comes to constancy and illusion, he argues that we cannot simply associate constancies with veridical perception and illusions with non-veridical perception. The relationship between veridical perception, constancies, and illusions is much more complex—for instance, non-veridical experiences need not be illusions and we can have constancy with non-veridical experiences. This complexity—and at times uncertainty—is due to issues at the core of our notion of veridical perception. Chief among them is that the crucial idea of correspondence with reality is imprecise.

⁹⁷ There is also the issue of which shape we are supposed to take a large three-dimensional object to be given the perspective from which we view it. For example, Schwartz considers seeing an elephant from up close, from behind, and from above (p. 392). Maund (2012) also discusses this issue in detail, noting how we can correctly perceive a table to be rectangular from different vantage points while noticing that its “apparent shape” may change slightly as we move. He argues that our veridical experience of a rectangular table can have illusory aspects because of how it appears from different perspectives, despite its unchanging intrinsic shape. Hatfield (2016) discusses this issue too. Maund links this ambiguity to the presence of different notions of constancy—one that stresses the intrinsic properties of objects; the other, tendencies for our experience to hew closer to intrinsic properties than retinal images. Like Schwartz, he argues that the relationship between veridicality, constancy, and illusion is complex. However, he focuses principally on issues that are distinct from what concerns us most in Schwartz’s discussion.

In order to determine if experience corresponds with reality there are three things we must get clear on. All three are alluded to in the long quotation above about the 250-foot-tall building. First, we need to know what physical reality is like. With respect to the low-level spatial properties of objects I will assume there is little controversy about physical reality. Second, we must determine what appears to us in experience in the sense of what our experience presents to us or purports to be about. For example, we may think my experience of a 250-foot-tall building is an experience where I am presented with a very tall building, or a building taller than my apartment building, or a building over 200 feet, but not one where I am presented with a building that is 250 feet tall—that is, we may not think that what appears to me is precisely a building that is 250 feet tall or that my experience encodes a precise height for the building. Third, we have to specify our standard of correctness. Two things can be said to correspond given a loose standard of correctness but fail to correspond once we make the standard stricter. Once we know what objects are external to us and have decided how to characterize the sensory properties that appear to us in experience, we still need to decide how strictly these two must correspond for experience to be considered veridical. Often our standard will be relatively, but not too, strict. For example, suppose I am looking at my table, chair, and ball, and the chair is exactly twice as far from the table as the ball is. Due to my perspective, I experience the chair as farther from the table than the ball, but perhaps not precisely twice as much. I doubt we would consider such an everyday experience to be non-veridical or to fail to correspond to how things are.⁹⁸

⁹⁸ Lack of clarity about veridicality and correspondence makes it hard to get clear on whether or not there is constancy or illusion. As Schwartz points out, claims that experiences exhibit constancy or are instances of illusion will often seem somewhat arbitrary or be based on expectations and prior experiences (see, e.g., p. 387, 395).

As Schwartz argues, our standard of correctness can vary depending on the context and the task at hand. Though he observes that our perception of spatial properties can be inaccurate in many respects, such failures of correspondence (given a strict standard) do not trouble us in daily life.

As discussed, we are not especially accurate assessing metric or comparative spatial properties as measured in the physical sciences. The fit is much better when vision serves to guide behavior in the nearby environment. Within limits, vision is reasonably successful informing us of what we can reach, lift, and manipulate. It is also reasonable successful in allowing us to avoid bumping into things that harm and guiding us to things that are further afield. When vision succeeds in these ways, I have no qualms saying experience corresponds to reality, indeed physical reality. What is accessible and useful at one time, for one person or animal, however, need not be the same at other times and for other perceivers. (p. 399-400)

We may not entirely agree with Schwartz here, or may find his attitude toward veridicality too liberal. However, the point remains that in order to determine if experience corresponds with reality the standard of correctness must be specified, and this is something done by us.

2.2. Characterizations of Veridical Perception in the Philosophical Literature

Before turning to mismatches and the question of whether or not we should consider most mismatches to be cases of veridical perception, I want to first look at how veridical perception is typically characterized in the philosophical literature.

In the philosophical literature on perception, illusions are typically characterized as cases where things appear to be a way they are not or instances when the way an object seems is different from how that object is (see, e.g., Robinson, 1994; Smith, 2002; Brown, 2008; Fish, 2009; Brewer, 2011; Macpherson, 2013; Macpherson and Batty, 2016; McLaughlin, 2016). Most often this is put in terms of properties, and an illusion is when an object *o* appears to be *F* when it is not actually *F*. In the case of vision, an illusion is when something looks *F* when it is not actually *F*. The disparity is between what appears to the individual—the sensory properties in experience—and the external object(s) present. For instance, the wall appears to be red when it is actually white, or we experience some object as being a square when it is round, or two things seem to be the same size when they actually differ in size. In cases of illusion we misperceive an object in a certain way, i.e. our perception of the object is not accurate in some respect. Our experience in such cases will often be accurate in many other respects—i.e. we may accurately perceive many other properties of the object in question—but it is illusory with respect to the property or properties *o* seems to have but which *o* does not actually have.

Since illusions are contrasted with veridical perceptions, this suggests that cases of veridical perception are cases where things appear to be the way they actually are. And when veridical perception is described this is what we find. Put in terms of objects and properties, veridical experiences are when an object *o* appears to be *F* when it actually is *F*. Furthermore, given that an experience will be illusory when only one property of an object is misperceived, veridical cases must be cases where no property is misperceived, i.e. where our experience is accurate in all respects.

This way of characterizing illusions and contrasting veridical experiences with illusions implies rather strict standards for veridical perception. The impression one gets from the

philosophical literature is that in order for an experience to be veridical it must be accurate in all respects, as any difference between what appears to us and the relevant external objects would render the experience an illusion. Typically the focus is on properties like colour, shape, and size. Thus, I have veridical experiences when I perceive the white wall as white, the round object as round, and the identical lines as equal in length.⁹⁹

But while strictness is implied by the claim that veridical perception is accurate perception and by the sorts of cases pointed to as examples of illusion, it is often not clear whether people are actually committed to a very strict correspondence or if they even need to be, given their other theoretical commitments. While discussions of perception and illusion typically revolve around the same sorts of properties and examples, two important issues are rarely, if ever, discussed. First, precisely which properties of an object we must accurately experience for that experience to count as veridical. Second, just how much our experience and the object perceived must agree. Take, for example, my experience of my computer and the width of its screen. Just how closely the shape or width I experience must correspond to the actual shape of the screen is unclear since accuracy is a matter of degree. If I draw two lines that are both 10.15 centimeters we may be perfectly happy to say they are the same length even if one is 10.1532 centimeters and the other is 10.1508 centimeters. In the same way, my experience can correspond to the actual shape of the screen even if it is not precisely accurate.

Those defending theories emphasizing the causal aspect of perception sometimes suggest that a looser correspondence between experience and external objects suffices for veridical perception. According to Sollberger (2015), we have veridical experience when a mind-

⁹⁹ I will focus on these sorts of more basic properties, rather than higher level properties like being a barn or a snake or a horse. Doing so is in line with my discussions in the previous chapters.

dependent idea is appropriately caused by certain mind-independent objects and preserves the higher-order relational and structural properties of those mind-independent objects. This sort of similarity between the properties present in experience and the properties of objects does not require that the two correspond precisely, but simply that there is the appropriate sort of mapping from one to the other. Similarly, Coates' (2007) causal view and navigational account of perception would seem to allow for experiences that do not strictly correspond to the mind-independent objects that cause them to count as veridical. And while he at times echoes the common refrain on illusions—i.e. as being cases where an object appears in a way that differs from how it is—he notes that “the line between normal perception and illusion is indeterminate” (p. 61).

While illusions are often taken to simply be cases where there is some difference between how things seem in experience and the relevant external objects, they are sometimes thought to consist in our falsely judging that some external object has a property it actually lacks. Naïve realists, in particular, often characterize illusions in this way (see, e.g., Brewer, 2011; Genone, 2014). On such approaches, illusions may very well be cases where *o* looks *F* even though it is not *F*, but the error here is in our judging that *o* is actually *F*, not in its looking *F*. This is because, as we have seen, some naïve realists argue that looks are objective properties of mind-independent objects: some object may look *F* even though it is not *F* because it has a certain look.¹⁰⁰ Whereas all illusions will be cases where *o* looks *F* even though it is not *F*, not all cases

¹⁰⁰ This focus on judgment when it comes to illusions is not exclusive to naïve realists, and we have already seen how some naïve realists opt for other approaches to illusions (Martin, 2004; Fish, 2009); but, it is clear why such an approach to illusions would appeal to the naïve realist. Naïve realists argue that in experience we are acquainted with mind-independent objects and their properties. Claiming that illusions are cases where we are acquainted with some property a mind-independent object actually has, but that we are simply misled into falsely judging that the object is some way it is not, allows naïve realists to understand illusions as object-involving. The

where o looks F even though it is not F will be illusions. According to Genone, “perceptual illusions can be understood as experiences that tend to produce false judgments” (2014, p. 340). Appearances can be misleading, but what appears to us is not itself mistaken or erroneous. Rather, “perceptual error belongs to the mistaken judgments one forms on the basis of [appearances]” (p. 353). While Genone and others may not explicitly claim that appearances are in a sense always veridical, that seems to be the view:

...there is no question of perceptual experience going wrong in the case of illusion. The mistaken beliefs we form on the basis of experience are the fault of judgment—taking the world to be a way that it is not. Perceptual experiences are misleading when we are insufficiently familiar with the environmental conditions surrounding the objects we are perceiving, such that their appearances lead us to draw false conclusions. (p. 363)¹⁰¹

At the very least, on such an approach, veridical perception can include experiences where what appears to us differs from how objects actually are, so long as such experiences do not mislead us and do not tend to produce false judgments.

I have already shown that such approaches to illusions cannot explain mismatches. I am not sure if such approaches offer an adequate account of the cases naïve realists typically take to be illusions—I am sceptical that they do, for reasons similar to those raised regarding mismatches, but the issue is beyond the scope of my project. However, characterizing illusions

error, according to such views, is in our failing to recognize that the appearance is merely a look and taking the object to actually be some way it is not.

¹⁰¹ See, also, Macpherson:

An alternative view of illusions, inspired by a disjunctivist view of perception, is that in illusions we have accurate experience but form inaccurate beliefs about the world on the basis of that experience. (2013, p. 6, fn. 6)

as experiences that may be misleading is attractive. Doing so allows for the possibility of things appearing to us a way they are not, yet this not being a case of illusion.¹⁰²

Although the contrast between veridical perception and illusion features centrally in debates about the nature of perception, it is frequently acknowledged that the line between veridical and non-veridical is unclear and that there will be hard cases about which it is difficult to know what should be said. It may be fair to be content with simply gesturing towards these issues if it is normally clear enough whether or not an experience is veridical and if illusions are relatively rare. But what the empirical work shows is that hard cases are common. Thus, we need to start thinking about them in more detail.

2.3. Mismatches as Veridical

Mismatches are cases where there is some difference between what is external to the individual and what they experience, between the properties of the relevant mind-independent objects and events and the properties that they are aware of in perception, the sensory qualities that appear to them. I have discussed how mismatches can arise due to binocular rivalry, multimodal perception, temporal filling-in, visual completion, normal geometrical misperception, colour perception, and speech perception. Sometimes these mismatches will be significant, as with the binocular rivalry study with patchwork images or in the McGurk effect. But in our daily lives mismatches will often be minor, involving only subtle discrepancies

¹⁰² While characterizing veridical and non-veridical perception in this way may be important for the naïve realist, such approaches to the veridical/non-veridical distinction in no way commits one to naïve realism. To illustrate, such an approach seems similar to Schwartz's, but his discussion is neutral between different theories of the nature of perception. Whereas the naïve realist will consider a case where *o* seems *F* even though it is not *F* to be veridical because they take us to perceive some mind-independent look of *o*, one may simply consider it veridical because we are not misled in thought or action.

between what our perceptual experience is like and the external things before us in the world. My focus here is on these minor mismatches that often occur in our day-to-day lives.

Are mismatches, however minor, instances of veridical perception? It would seem like they are not. Based on the common idea of veridical perception as instances when we perceive things as they actually are, mismatches are by definition not veridical. Mismatches are experiences that are not completely accurate.

But that we get on in the world very well suggests that most of our perceptual experience is perfectly good. At the start of section 2 I noted that for the most part our perceptual experience seems to us to be accurate. Given the ubiquity of mismatches, many experiences we would typically be perfectly happy to consider good cases of experience will involve minor mismatches. I do not think that we should relinquish the idea that most perception is veridical perception. Given that our perceptual experience helps us get on in the world just fine and mismatches occur in perfectly normal settings, our notion of veridical perception should be such that lots of minor mismatches count as veridical.

The fact that we can adopt varying standards of correctness gives us the opportunity to consider most mismatches as veridical perception—it gives us the room to do so. As Schwartz points out, when it comes to the question of whether an experience corresponds with reality we must determine how strict of a correspondence is required in order to be considered accurate or correct. While all mismatches are inaccurate in some way, these experiences are also accurate in many respects. And when it comes to assessing the accuracy of an experience, we can consider how closely it corresponds to the external world in different ways. We may look at how many properties the experience gets right, i.e. regarding which properties there are discrepancies versus those where there are not. Or, we may look at how closely the experience must correspond with

the external objects regarding a particular property, i.e. even if there is some discrepancy regarding a property, we may decide it is minor and falls within the acceptable margin of error so the experience should count as corresponding to the world. We must keep both of these dimensions in mind when considering how strict our standards of correctness should be—our experiences typically do not involve only a single object or property. If we are being strict, mismatches are not completely accurate experiences. But why adopt such strict standards? Why think experiences must be completely accurate in order to count as veridical?

We may be tempted to say that minor mismatches are non-veridical for a few different reasons. However, as we will see, each is problematic. Note that one cannot simply say that mismatches are non-veridical because there is a lack of correspondence. The issue here is how strict the required correspondence should be, or why we should have a strict notion.

One reason why we may want to say that minor mismatches are non-veridical is if they are due to our perceptual system malfunctioning in some way. If an experience is inaccurate because one's perceptual system fails to work in the normal way, we may think such an experience should be considered non-veridical. Such a failure could be due to some sort of individual deficit or an otherwise normal individual being placed in a situation where the perceptual conditions are abnormal. We may think that certain laboratory conditions are abnormal because experimenters contrive to trick our perceptual systems in some way. This could be true about binocular rivalry studies and cases like McGurk studies where auditory and visual stimuli are manipulated in a way we are unlikely to see in everyday life. However, while experimental conditions are artificial, they are often created expressly to mimic certain aspects of the normal world. Furthermore, the sorts of minor mismatches I am suggesting should be considered cases of veridical perception occur in our day-to-day lives, in ecological settings.

Such mismatches arise in perfectly normal conditions, in normal individuals—in fact, they result from the normal workings of our perceptual systems. Thus, while the result of our perceptual processing may be a mismatch, this is not due to the situation being abnormal or our perceptual system malfunctioning in some way.

A second reason we may have for considering mismatches to be non-veridical is if they misguide us behaviourally. We may want to say that a mismatch cannot be a case of veridical perception if it leads to some sort of behavioural blunder, e.g. our bumping into something or failing to grip a nearby object. The idea, here, is that the very discrepancy in virtue of which the experience is a mismatch leads to the behavioural error. Such a principle seems plausible to me. However, most mismatches will not lead to any sort of behavioural error. In most cases, the mismatch element will not concern our action, in particular not at that moment. For example, consider walking down the street. Though our actions are guided by our experience, there are lots of things that will be present in our experience that are irrelevant to our behaviour here. And even when it comes to spatial properties we may misperceive, our experience of these things will change as we approach the objects or as our perspective changes. Despite mismatches being ubiquitous, things do not often go awry behaviourally due to such discrepancies.¹⁰³

However, we may think that whether or not a mismatch actually leads to such an error is irrelevant. That is, we may think that what is important is that an experience has the potential to misguide us behaviourally. If so, even when we do not act on the basis of the mismatch element, we should consider mismatches to be non-veridical because they have the potential to mislead us in this way. I do not find this idea as appealing as the previous one. But, in any case, most

¹⁰³ Though Schwartz makes this point—quoted above—in the context of constancies and the sorts of inaccuracies he discusses, it is no less true that in general vision guides us quite well in our nearby environment.

mismatches do not even have the potential to misguide us behaviourally. This is because the minor discrepancies will often not involve things that will impact our actions and behavioural decisions. For one, any perceptual experience where we are relatively stationary will not have the potential to misguide us. More importantly, as noted in the previous paragraph, lots of things we hear or see (or taste or smell or touch) are irrelevant to our behaviour. As I walk down the street, many of the sensory qualities I am aware of do not in any way impact my behaviour. And even if we consider only vision and spatial properties, lots of things I see—whether in the distance or on either side as I pass—are unrelated to the actions I perform in such a way that a mismatch does not even have the potential to misguide me. If my momentary experience of some buildings in the distance is slightly inaccurate due to some visual completion or the sorts of things discussed by McLaughlin, it is hard to see how this could even potentially misguide an action.

Finally, we may think that we should consider mismatches to be non-veridical if they mislead us when it comes to belief.¹⁰⁴ Once again, we can ask whether the experience would have to actually lead to some false belief or if it simply must have the potential to do so. An experience having the potential to lead to a false judgment does not seem like a good reason to consider it to be non-veridical. How large or small things seem when we see them from a distance can lead us to falsely believe that some object is larger than another object it actually is smaller than—for example, we may think stars are smaller than the moon. But we do not typically consider such experiences to be non-veridical. While such experiences can mislead us, we are typically not misled because we understand that things look smaller to us when they are farther away.

¹⁰⁴ Insofar as our beliefs may guide us behaviorally this is not completely independent from the previous issue. However, there are also many cases where we make judgments that do not or could not possibly concern behaviour.

But the idea that an experience must actually lead to a false belief in order to be considered non-veridical is equally problematic. Many experiences we consider illusory do not actually lead to false beliefs because we are familiar with the effect and know they are illusions. These cases can, however, be tricky. On the one hand the lines in the Müller-Lyer illusion look to be unequal in length; on the other hand, we know that they are actually the same length. Similarly, if we are familiar with the McGurk effect we will not think that the speaker said [da], even though that is how it seemed to us. And, in any case, most mismatches will not mislead us because the mismatch element will be irrelevant to our goals or we will not care much about whatever element or property our experience has actually got slightly wrong. We do not form beliefs about everything in each of our experiences. As we get on in the world and perform various tasks there is lots we do not really notice or dwell on. An experience that is slightly inaccurate will tend not to mislead us if the way in which it is inaccurate is not terribly important to us.

Once we move away from the idea that veridical perception must be perfectly accurate it becomes difficult to say precisely when an experience is accurate enough to count as veridical. Whether or not we are misled in thought or action is an important consideration. But I am hesitant to argue that an experience is veridical so long as our perceptual system is functioning normally and the experience does not mislead us. Perception can help us get on in the world even if it is quite inaccurate in some respects, and there is the temptation to view an experience as in some sense defective if it involves significant mismatches, even when it reliably guides our actions.

We may want to say that an experience is veridical when there is a close enough correspondence between the experience and the mind-independent objects and events that cause

it and the individual is not misled. Or, perhaps, when the experience does not have the potential to mislead in a significant way. Determining whether the correspondence is close enough can depend a number of factors. We can consider whether the experience is completely accurate regarding most features, i.e. if in general the experience corresponds with the external world. We can also consider how significant the discrepancies are between the experience and the external world regarding the features where there are mismatches. Another consideration may be the individual's goals or what they are doing while having the experience. We may have good reason to consider an experience veridical in one context, given the tasks or goals of the individual at that time; but, we may consider the very same experience non-veridical if it was being put to a different use. There is no need for our standards of correctness to be unchanging given the variety of uses to which we put experience and the various contexts perceivers find themselves in.¹⁰⁵

An analogy with communication and language may be helpful. We may think there is successful communication when things go smoothly and there is a happy outcome, when a person is more or less understood to the degree necessary for the task at hand. But we may also

¹⁰⁵ The point here is not about what perception is for in general, but what we are doing in a particular instance. See Schwartz (e.g. p. 400) for further discussion of how our needs may vary and can be met in different ways by perceptual experiences. However, in considering the issue of veridicality we may also want to reflect on what perception is for more generally. For example, in discussing visual constancies Hatfield (2016) argues that:

...perceived shapes, sizes, and distances can be appropriate for guiding behavior even if they do not fully match physical shapes, sizes, and distances. ...Accordingly, we should not assume that the aim of the visual system is to recover and phenomenally represent properties such as shape and distance as they are physically. Rather, it will be enough if the sizes, shapes, and distances presented phenomenally are systematically related to physical values, sustain action guidance, and can support cognitive discernment of physical values. (p. 152)

While it is true that mismatches will differ from the sorts of things Hatfield describes in that such experiences will not be so systematically related to physical objects, the experiences may nevertheless succeed in serving the same purposes, despite this.

have a stricter conception of successful communication. If we think that what speakers mean are propositions, we may think that communication succeeds only when the hearer picks up the very proposition the speaker meant. We can harbour both views of communication simultaneously, and when confronted with particular cases find ourselves pulled in different directions when trying to judge whether communication is successful.¹⁰⁶ But if what speakers mean is often indeterminate—in the sense that there is no particular proposition meant—we may have to give up the stricter notion of communication, since many seemingly good cases of communication will not involve—indeed, could not possibly involve—such a correspondence (Buchanan, 2010, 2013; Sperber and Wilson, 2015; Neale, 2017). Instead, we consider communication to be successful when what is meant and what is picked up is related in the right sort of way, or if there is enough overlap between the two. How much overlap is required could depend on the context and goals of the communicative exchange (Sperber and Wilson, 2015).

In language, the indeterminacy of meaning pushes us towards a conception of communication where communication can be successful even if what the speaker means and the hearer picks up do not correspond precisely. We may still feel the pull of the stricter conception, but it makes no sense to say that communication is rarely successful because meaning is often indeterminate. In perception, the ubiquity of mismatches pushes us towards a looser notion of veridical perception, where an experience can be veridical even if it does not correspond precisely with the external world. We may still feel the pull of the idea that an experience must be completely accurate in order to count as veridical, but it seems odd to say that perception is

¹⁰⁶ If we have studied philosophy of language, we may find ourselves more inclined towards the stricter notion, more compelled to require that what the speaker means and hearer picks up correspond precisely in order for communication to be successful.

rarely veridical given that perception helps us get on in the world quite well and most mismatches are minor.

Indeed, one of the lessons from Schwartz's discussion of constancies is that it is not so clear what it means to perceive things as they really are (see also Maund, 2012; Hatfield, 2016). The idea of an experience corresponding with the external world is not so straightforward. There is certainly an important difference between constancies and mismatches: mismatches are cases where there is a clear difference between our experience and the intrinsic properties of the external objects before us. But, whether we consider the issue of veridicality in the context of constancies or mismatches, the fact remains that there is work to do in order to decide the standards of correctness. As Schwartz notes:

We set the standards of veridicality, not the world. Although influenced by history and causal connections, these facts do not and cannot fix perceptual norms. We gerrymander standards for the varied purposes at hand. ...There is no privileged physical, or for that matter phenomenal specification, of perception's goal(s), and no unique standard for assessing appearance correctness and error. (p. 401)

More needs to be done in order to articulate an appropriate notion of veridical perception, and even once we have done so it may continue to be difficult to know what to say about certain cases. But our notion of veridical perception should be one on which most everyday mismatches are instances of veridical perception.

3. Predictive Coding and Indirect Realism

I have noted on a few occasions, while discussing perceptual processing, that the brain must make sense of its sensory input or that we may understand the brain as computing what is

most probably out there, external to us. There are different theories about such neural processing, different views about the computations involved in perceptual processing. We may even understand these different theories as different accounts of how we solve things like the inverse optics problem McLaughlin discusses (see 2.5 of previous chapter). A theory that currently enjoys quite a bit of popularity is the predictive coding view, according to which the brain is an organ of prediction (Friston, 2010; Hohwy, 2013, 2016; Clark, 2013a, 2013b, 2014, 2017). According to this theory, “the brain is a sophisticated hypothesis-testing mechanism, which is constantly involved in minimizing the error of its predictions of the sensory input it receives from the world” (Hohwy, 2013, p. 1).

The predictive coding view is offered as a unified theory of the brain that can explain perception, action, and cognition. Whether the theory suggests a certain view about our relationship to mind-independent objects in perception has been the subject of recent discussion. Hohwy argues that the theory entails that perception is indirect in the sense I have been concerned with. And, more generally, the theory has served to motivate claims that recent scientific work in perception supports indirect realism (see, e.g., Frith, 2007; Metzinger, 2009; Westerhoff, 2016). Here, I look at whether Hohwy is right to think that if predictive coding is correct, then perception is indirect. I will not assess the merits of the theory itself or consider the question of whether or not we should take predictive coding to be a correct theory of the brain. Rather, my focus is simply on whether the theory entails indirectness, as Hohwy argues.¹⁰⁷

The predictive coding view takes the brain to be principally concerned with minimizing the difference between its own predictions about the way the world is and the sensory input it receives. According to the view, the brain generates models or hypotheses that predict its sensory

¹⁰⁷ Evaluating the theory is beyond the scope of my project. While the theory has a number of proponents, it also faces substantial opposition (see, e.g., replies in Clark, 2013a).

input, which it then can test against its actual sensory input. The brain strives to reduce the difference between its predictions and input, or to minimize the error in its predictions. Since the brain is viewed as a biological machine for minimizing prediction error, the view is sometimes called “prediction error minimization theory.” And because the brain’s hypotheses are understood in terms of Bayesian inference, or within the Bayesian framework, such approaches to perception may be termed “Bayesian.” While I will focus on perception, predictive coding is offered as a rather ambitious view about the brain and mind:

PEM says that prediction error minimization is the only principle for the activity of the brain. This means that the brain is an organ that on average and over time continually minimizes the error between the sensory input it predicts on the basis of its model of the world and the actual sensory input.

... If this is all the brain does, then perception, action, attention, and all other mental processes, must come down to prediction error minimization. (Hohwy, 2016, p. 260)

[The theory] offers a distinctive account of neural representation, neural computation, and the representation relation itself. It depicts perception, cognition, and action as profoundly unified and, in important respects, continuous. And it offers a neurally plausible and computationally tractable gloss on the claim that the brain performs some form of Bayesian inference. (Clark, 2013a, p. 187)

Hohwy takes the predictive coding view of the brain to entail the indirectness of perception. His understanding of perception starts with sensory input. Mind-independent objects and events interact with our bodies, causing sensory input, which the brain receives. The brain then works to determine the hidden causes of the sensory input. Since the same input can result

from a variety of different causes, this is no simple task. On the predictive coding view, the brain uses Bayesian inference to generate hypotheses about the causes of its sensory input.

This picture casts the brain as self-evidencing. The hypotheses it generates about the causes of its sensory input are judged against that sensory input. A model of the brain is successful insofar as it explains away the very evidence—the sensory input—that supports the model. As Hohwy notes:

When h_i is self-evidencing, there is an explanatory-evidentiary circle (EE-circle) where h_i explains e_i and e_i in turn is evidence for h_i

The internal model that generates hypotheses that over time makes the evidence most likely, and does so most precisely and simply, will have its own evidence maximized. That is, as a model generates hypotheses that explain away occurring surprising evidence (i.e. minimize prediction error) it maximizes the evidence for itself. Prediction error minimization thus constitutes self-evidencing. (2016, p. 263-4)

The brain is self-evidencing because there is an evidentiary boundary: on one side of the boundary are the sensory input and the brain's models, and on the other side are the hidden causes of that input. The causes that are on the other side of the evidentiary boundary can only be inferred.

In an ideal but impossible design, perception, attention and action would require the brain to simultaneously access both the internal estimates and the true states of affairs in the world. This would allow it to compare the representation and the represented, the attended and what is worth attending to, action planning and what is acted upon. Philosophers have long recognized that there is no such access since we never have unfettered knowledge of states of affairs in the world. PEM delivers the

tools for circumventing this problem: it gives the brain access to two things it can compare, namely the predicted and the actual input. (p. 262)

Perception thus involves the brain inferring the causes of its sensory input, and the computations involved here are Bayesian. The idea is that the brain implements Bayesian rules, where prior knowledge or assumptions lead to hypotheses formed on the basis of probability that are then tested against the evidence—i.e. the sensory input—resulting in a posterior probability for a given hypothesis.¹⁰⁸

Perception is indirect because, according to Hohwy, our experience is determined by our internal models. Hohwy characterizes Bayesian inference as active inference on the basis of internal models that represent the world. The bottom-up sensory signal serves as feedback to the top-down predictions, or feedback to the internal models of the world. But perceptual experience depends on this top-down processing and is not primarily driven by the sensory input and bottom-up signals. The character of our experience is thus determined by the inner models, rather than the mind-independent objects that cause our sensory input. As Hohwy claims, perception “is indirect in the sense that what you experience *now* is given in your top-down prediction of your ongoing sensory input, rather than the bottom-up signal from the states of affairs themselves” (2013, p. 48).¹⁰⁹ This is just the sort of indirectness I articulated in the first chapter. According to Hohwy’s presentation of the predictive coding view, we are causally related to mind-independent objects and events, but while our internal models represent the world they stand at the opposite

¹⁰⁸ For more on Bayesian inference in the context of perception, see Hohwy (2013).

¹⁰⁹ He expresses this idea in various ways: “The hypothesis that is selected determines perceptual content...” (p. 37); “On our account, this is captured by the idea that perceptual content is maintained in internal models...” (p. 227); “This model is personal-level in the sense that it determines perceptual content...” (2016, p. 267); “What ultimately determines the resulting conscious perception is the best hypothesis...” (Hohwy et al., 2008, p. 690).

end of the causal sequence and can exist independently of the things that cause them and which they represent. The sort of directness essential to naïve realism, where mind-independent objects are constituents of our experiences, is therefore ruled out. Perception involves our mentally representing the world, such that the character of our experience is determined by an internal representation.¹¹⁰

It is clear that the evidentiary boundary is pivotal for Hohwy's understanding of perception. By taking sensory input as the starting point, the brain and mind is in an important sense isolated from the mind-independent causes of those inputs (he calls this "seclusion"). But there are two related questions we should ask. First, how is this boundary related to the predictive coding and Bayesian elements of the theory? And, second, is this boundary unique to the predictive coding view?

Hohwy appears to offer prediction error minimization as a way we overcome the evidentiary boundary, indicating that the boundary can be understood independently from

¹¹⁰ Although Clark appears to take issue with Hohwy's arguing that predictive coding entails indirectness, it is not clear how much he actually disagrees with the way Hohwy understands perception. On the one hand, he is critical of Hohwy's emphasis on seclusion and notes:

The intervening mechanisms thus introduce no worrisome barrier between mind and world. Rather, it is only *because* of such sub-personal complexities that agents like us can be perceptually open to the world itself. (2013a, p. 199)

However, he (2014) also claims that predictive coding poses problems for disjunctivism and seems to side with Hohwy when saying:

There remains a genuine sense in which the *experienced world* may be said to be constructed by the brain from behind an evidentiary boundary imposed (currently at least) by the biological senses. (2017, p. 735)

Part of the dispute seems to hinge on the use of "perceive" and Clark's stressing that we perceive the external world, whereas Hohwy at times claims that "[w]hat we perceive is the brain's best hypothesis" (2007, p. 323) or that it is "difficult to say that what we perceive is in some sense directly or precisely the world" (2016, p. 267). However, elsewhere Hohwy claims that we perceive the external world (e.g. 2013, p. 13) and—as I noted in the first chapter—there is no actual tension between indirectness and our perceiving external objects. But the main issue Clark seems to have with Hohwy's presentation of the predictive coding view concerns embodiment, and not indirectness. It is Hohwy's emphasis on the brain—by itself—being isolated from the world that Clark finds objectionable (see, in particular, Clark, 2017).

prediction error minimization, or as an independent commitment. That is, while the boundary may be part of the more general view, it is independent from—and prior to—a commitment to Bayesian processing. He defines his own “problem of perception” as the problem of how the brain manages to perceive the external world on the basis of sensory input, or how the brain manages to figure out the causes of its sensory input. “The problem of perception is a *problem* because it is not easy to reason from only the known effects back to their hidden causes” (2013, p. 13). Prediction error minimization is then offered as the mechanism whereby we may infer those hidden causes and solve the problem:

Unconscious Bayesian perceptual inference is an attractive way to begin to deal with the problem of perception. Inference is needed to overcome the brain’s encapsulation in the skull... (p. 41)

But while predictive coding is one theory of how the brain may try to figure out the causes of its sensory input, it is far from the only one (see, e.g., Purves and Lotto, 2003; see also discussion in McLaughlin, 2016). Rather than being a unique feature of the predictive coding view, the boundary seems to be part of the more general scientific causal picture of perception. According to this picture, causal interactions between our sense organs and the external world yield sensory input, which initiate a series of stages of processing in the brain, eventually resulting in perceptual experience. Although many argue that this picture suggests that perceptual experience involves inner states that are distinct from the external objects that cause them, naïve realists contend that this is mistaken and that they can accept that experiences are part of the causal order. Naïve realists will accept the causal picture yet argue that we should understand perception as a whole to involve our direct awareness of mind-independent objects (Le Morvan, 2004; Martin, 2004; Nudds, 2009). It is not clear if naïve realists are right to think their theory is consistent with the scientific causal picture of perception. But, it is at least not obvious that the

sort of boundary Hohwy discusses entails indirectness, as it is a part of other views that do not clearly necessitate indirectness.¹¹¹

It is certainly true that the picture Hohwy offers is similar to the one I defend in the previous chapter (in section 5), which begins with sensory inputs and ends with the brain constructing a representation. (Whereas Hohwy and company argue that the intervening processing is Bayesian, I do not do so.) However, in my case, I initially argue that we should take the brain to construct a representation because this is how we can best explain our experience in mismatch cases. It is not on the basis of an alleged boundary that I argue we should take perception to involve mental representation and be indirect.¹¹² So, does Hohwy do something analogous and give us good reason to think predictive coding entails indirectness?

An idea that seems to be popular nowadays is that if perception is active, can be understood as construction, or involves top-down processes, then it must be indirect. Hohwy often contrasts his top-down view with bottom-up approaches to perception. But there are different theories about such top-down, feedback, or active processes in the brain, aside from predictive coding. If such processes are the reason we are supposed to think of perception as indirect, then this would seem to not have anything to do with predictive coding per se or whether we should understand such processing as Bayesian. Moreover, it is not entirely clear why such processes are taken to render perception indirect.

¹¹¹ Orlandi (2016) disputes Hohwy's interpretation of the predictive coding view. She argues that Hohwy and others in the predictive coding camp downplay the role of the sensory input and bottom-up processing.

¹¹² Although I am sympathetic to the idea that the scientific causal picture suggests indirectness, I think that more needs to be said to support this stance. In my case, this involved arguing that the naïve realist could not explain our experience in mismatch cases. Another approach would be to explain why naïve realism is incompatible with the idea that perception is causal, as Coates (2007) does; for a reply, see Hobson (2013).

Something like the mismatch argument may be suggested by Hohwy. He points to binocular rivalry and how it demonstrates the role of the brain in determining what we experience. Binocular rivalry, he says,

...puts pressure on the idea that perception is purely stimulus driven, bottom-up feature detection. During rivalry, the physical stimulus in the world stays the same and yet perception alternates, so the stimulus itself cannot be what drives perception. (2013, p. 20)

And he certainly thinks that the brain always functions in the same way.¹¹³ But if this is how Hohwy is reasoning, then what leads us to think that perception is indirect is not the active processes themselves, but that these processes lead to a mismatch in the rivalry case—or, at least, that the rivalry case makes evident that the brain, and not the external stimulus, determines the character of our experience. However, that such top-down or active processes are always involved in perception does not mean that mismatches always result. Although I argue that mismatches are the norm, this does not simply follow from the fact that perception involves active processes. And, on the mismatch argument, we generalize to all experiences not simply because there are active processes at work, but because a representation is constructed by the brain. There may be an argument from the mere presence of such top-down or active processes to indirectness, but it is not quite clear what the reasoning is—and, an explanation of the tension between such processes and directness is not to be found in Hohwy's work.

It is therefore not clear whether predictive coding entails that perception is indirect. Even if the theory does have this consequence, this would seem to be due to the general picture of the

¹¹³ For instance, while discussing rivalry he notes:

It seems reasonable to work on the assumption that the brain always uses some sort of inferential process to perceive the world, and that rivalry is an effect that simply makes the brain's everyday inferential processes easier to spot. (p. 20)

brain and world Hohwy takes as a starting point, rather than anything about prediction and Bayesian inference—i.e. not because of the very features that distinguish the theory from competing theories of perceptual processing. If it is the case that the internal models of predictive coding theory must be understood as representations (Gładziejewski, 2016), then we may think this also suggests indirectness. But, first, whether these internal models should be understood as representations is a point of contention (Orlandi, 2016). And, second, even if they are to be understood in this way, unless these models determine the character of our experience, the predictive coding theorist's case for indirectness would be similar to Nanay's, discussed in the second chapter, in which case the direct theorist has a reply. I have argued, anyway, that perception is indirect. And, if predictive coding is correct, then whether or not the theory itself entails indirectness may be unimportant. However, I am not sure that indirectness follows simply from the predictive coding view.

4. Perception, Indirectness, and Interpretation

In this dissertation I have argued that perception involves our mentally representing the world, and the phenomenal character of our perceptual experience—the sensory qualities that appear to us in experience—is determined by an inner representation. I defended two arguments rooted in recent empirical work on perception. These arguments show that perception is, in an important sense, indirect, and that naïve realism is false.

In the first chapter I argued that whether or not perception involves mental representation is a central question. Views that take what appears to us in experience—the phenomenal character of experience—to depend on representational mental states can be understood as indirect. On such views, mind-independent objects play a causal role in perception but are not

constituents of our experiences. This is the very sort of indirectness characteristic of views that accept the common kind assumption; such non-disjunctive views deny that mind-independent objects can have the sort of role in perception they do according to naïve realism, which takes us to be directly related to objects such that they are constituents of our experiences. While discussions of the common kind assumption feature prominently in recent philosophical debates about the nature of perception, I argued that focusing on the question of whether or not perception involves mentally representing the world was advantageous. This question places naïve realism in opposition to a family of views that includes indirect realism, critical realism, and intentionalism. I then argued that Sollberger's (2015) indirect realism can be understood as simply accepting that perception involves mental representation, and that critical realism and intentionalism each involve additional commitments. Thus, we can understand the question of whether or not perception involves mental representation in terms of the opposition between naïve realism and indirect realism.

In the second chapter I examined arguments that appeal to recent empirical findings in order to argue that naïve realism is false or that indirect realism is correct. I argued that there was one promising argument in this literature: the mismatch argument, presented by Smythies and Ramachandran (1997) in support of indirect realism. Smythies and Ramachandran argue that certain findings—regarding a binocular rivalry study with patchwork images and various kinds of filling-in—indicate that the phenomenal character of our experience is determined by a representation generated by the brain, rather than by external objects. Although I argued that it is not clear whether Brown (2008) also offers a version of the mismatch argument, I noted how some of the findings he points to can feature in the argument, which gets off the ground with

cases where there is a mismatch between what one's perceptual experience is like and the actual physical objects before us in the world.

I began the third chapter by taking a close look at empirical work relating to mismatches. I examined findings concerning binocular rivalry, multimodal perception, temporal filling-in, visual completion, spatial perception, colour perception, and speech perception. I argued that these findings indicate that mismatches—often slight, unnoticed, and inconsequential—are ubiquitous in day-to-day life. The ubiquity of mismatches led me to formulate a second argument in support of indirect realism: the argument that mismatches are ubiquitous.

I then defended the key claims found in the mismatch argument and the argument that mismatches are ubiquitous. First, I argued that the best explanation for our experiences in mismatch cases invokes mental representations, and that naïve realists are unable to adequately explain such cases in a manner consistent with their theory. Second, I argued that perceptual processing always operates in fundamentally the same way, so the phenomenal character of our experience is always determined by a representation constructed by the brain. Third, I considered what naïve realists might say in response to the ubiquity of mismatches and argued that they cannot rescue their theory by claiming that the naïve realist analysis applies only to non-mismatch cases. Thus, even without generalizing from mismatches to all cases of experience we are entitled to reject naïve realism.

Finally, in this chapter, I briefly explored how we may think about perceptual veridicality in light of the fact that mismatches are the norm. I argued that we should adopt a looser notion of veridical perception, allowing us to retain the idea that most perception is good, veridical perception.

A certain picture of perception emerges from the empirical work I have examined in this dissertation. Mind-independent objects and events causally interact with our bodies, generating sensory input, which initiates a series of stages of processing in the brain. This perceptual processing involves active processes—such as when the brain fills in the blind spot—as well as feedback from later to earlier stages of processing and interactions between different modalities and systems. In the later stages of this processing the brain generates, or constructs, a representation of the external scene. And it is this representation of the external world, reflected in the neural activity of the relevant brain areas, that determines what our perceptual experience is like, i.e. the sensory qualities that appear to us in experience.

An important consequence of the sorts of processes involved in perception is that mismatches commonly result from the normal functioning of our perceptual systems. While active and top-down processes help us process ambiguous and noisy sensory input in such a way that we can very successfully get on in the world, they also lead to often minor discrepancies between our experiences and the mind-independent objects and events that cause them.

This picture of perception is incompatible with naïve realism. Thus, the empirical work gives us good reason to reject naïve realism. Any theory that denies that perception involves mental representation or on which perception is taken to be direct in the way characteristic of naïve realism will be at odds with the empirical work. If such a theory is offered as a theory about the nature of perception—rather than, say, as an account of how things seem to us introspectively—then we have good reason to cast it aside.

Compatibility with empirical work need not be decisive when it comes to our philosophical theorizing. But, in the absence of considerations otherwise, it seems reasonable for empirical work to weigh heavily when it comes to matters about which empirical findings can

plausibly inform us. Whereas indirect realist approaches, under the banner of sense-data theory, have been unfairly maligned as non-naturalistic or incompatible with a scientific understanding of the mind (Martin, 2001; Crane, 2005; Crane and French, 2015), it is naïve realism that runs afoul of our empirically-grounded understanding of perception, despite its pretensions otherwise (Martin, 2004; Fish, 2009; Brewer, 2011; Phillips, 2016).

Though I cast things in terms of indirect realism versus naïve realism, indirect realism is not alone in taking perception to involve mental representation. Critical realism and intentionalism also accept the conclusion I defended—thus, my arguments would appear to support these views as well.

However, Masrour (2017) argues that some of the empirical findings I discuss pose a problem for certain intentionalist theories. In presenting McLaughlin's work on normal geometrical misperception, I discussed how angles may appear distorted when seen in different contexts or from different perspectives. Masrour points to such phenomena and notes how the same angle may look very different to us depending on these factors; therefore, two experiences of the same spatial property can be seemingly incompatible to us. Although he focuses on angles, Masrour rightly points out that our normal perception of other spatial properties may result in similar "visual dissonance." We can understand the apparent incompatibility in terms of the two experiences of the same angle having different phenomenal character. Masrour argues that intentionalist theories with a certain naturalistic view of content cannot account for the difference in phenomenal character of these two experiences. The empirical findings pose a problem for such views because of how they understand content, combined with the intentionalist claim that the phenomenal character of an experience is completely determined by its intentional content.

Masrou's argument only targets intentionalist theories with a certain view of intentional content, so intentionalists that opt for alternative accounts of content can avoid the objection. Whether or not the intentional properties of an experience completely determine its phenomenal character does not seem to be an empirical question—at least when considered abstractly in this way. However, as Masrou argues, empirical work may help us resolve family disputes once we consider the particular commitments of certain intentionalist theories, such as those resulting from the view of content adopted. In this way, empirical work may help shed light on the dispute between indirect realism and intentionalism if certain findings pose problems for the distinctive intentionalist claim about phenomenal properties.

The empirical work discussed in this dissertation would appear not to bear on issues related to conceptualization, leaving untouched any disputes between indirect realism and critical realism. However, we may wonder how significant the difference actually is between these two approaches, or whether the sort of indirect realism Sollberger defends is genuinely opposed to the conceptual component of critical realism.

Undoubtedly, more work is needed in order to flesh out our understanding of the nature of perception. Empirical work may be pivotal here, and certain findings may not only help to settle disputes between different theories but to deepen our understanding of the relationship between the brain and mind. Besides looking to empirical work, theoretical issues concerning views that feature mental representation deserve especially close attention. This may include closer scrutiny of the differences between indirect realism and critical realism, as well as examining the sorts of issues possibly facing intentionalism which I noted in the first chapter.

Perception is, in an important sense, indirect. This does not mean that we perceive some mind-dependent entities. Rather, we perceive mind-independent objects and events, which we

are causally related to. As philosophers have long pointed out, our understanding of perception is closely related to a variety of philosophical questions in philosophy of mind more generally, epistemology, and metaphysics. The epistemological and metaphysical consequences of various theories of perception have been the objects of particular concern. The precise metaphysical implications of indirectness—or exactly how we ought to interpret such a theory—is an interesting question (see, e.g., Westerhoff, 2016). I see no reason to be concerned about the implications of the indirectness defended here; however, these more speculative concerns lie beyond the scope of my project.

Bibliography

- Allred, Sarah and Jonathan Flombaum (2014). "Relating Color Working Memory and Color Perception," *Trends in Cognitive Science* 18: 562-5.
- Allred, Sarah and Maria Olkkonen (2013). "The Effect of Background and Illumination on Color Identification of Real, 3D Objects," *Frontiers in Psychology* 4.
- Allred, Sarah and Maria Olkkonen (2015). "The Effect of Memory and Context Changes on Color Matches to Real Objects," *Attention, Perception, & Psychophysics* 77: 1608-24.
- Anaya, Alfonso and Sam Clarke (2017). "Naïve Realism and Unconscious Perception: A Reply to Berger and Nanay," *Analysis* 77: 267-73.
- Appelbaum, Irene (1998). "Fodor, Modularity, and Speech Perception," *Philosophical Psychology* 11: 317-30.
- Appelbaum, Irene (1999). "The Dogma of Isomorphism: A Case Study From Speech Perception," *Philosophy of Science* 66: S250-9.
- Austin, John (1962). *Sense and Sensibilia*, Oxford University Press.
- Ayer, Alfred (1940/1963). *The Foundations of Empirical Knowledge*, MacMillan and Company.
- Berger, Jacob and Bence Nanay (2016). "Relationalism and Unconscious Perception," *Analysis* 76: 426-33.
- Bertelson, Paul and Béatrice de Gelder (2004). "The Psychology of Multimodal Perception," in *Crossmodal Space and Crossmodal Attention*, edited by Charles Spence and Jon Driver, Oxford University Press.
- Block, Ned (2004). "Mental Paint," in *Reflections and Replies: Essays on the Philosophy of Tyler Burge*, edited by Martin Hahn and Bjorn Ramberg, MIT Press.
- Block, Ned (2014). "Seeing-As in the Light of Vision Science," *Philosophy and Phenomenological Research* 89: 560-72.
- Bolognini, Nadia, Francesca Frassinetti, Andrea Serino, and Elisabetta Làdavas (2005). "'Acoustical Vision' of Below Threshold Stimuli: Interaction Among Spatially Converging Audiovisual Inputs," *Experimental Brain Research* 160: 273-82.
- Brainard, David, and Ana Radonjic (2014). "Color Constancy," in *The New Visual Neurosciences*, edited by John Werner and Leo Chalupa, MIT Press.
- Brewer, Bill (2007). "Perception and Its Objects," *Philosophical Studies* 132: 87-97.

- Brewer, Bill (2008). "How to Account for Illusion," in *Disjunctivism: Perception, Action, Knowledge*, edited by Adrian Haddock and Fiona Macpherson, Oxford University Press.
- Brewer, Bill (2011). *Perception and Its Objects*, Oxford University Press.
- Brewer, Bill (2017). "The Object View of Perception," *Topoi* 36: 215-27.
- Broad, Charlie (1923). *Scientific Thought*, Kegan Paul, Trench, Trubner & Co.
- Brouwer, Gijs and David Heeger (2013). "Categorical Clustering of the Neural Representation of Color," *Journal of Neuroscience* 33: 15454-65.
- Brown, Harold (2008). "The Case for Indirect Realism," in *The Case for Qualia*, edited by Edmond Wright, MIT Press.
- Buchanan, Ray (2010). "A Puzzle About Meaning and Communication," *Noûs* 44: 340-71.
- Buchanan, Ray (2013). "Conversational Implicature, Communicative Intentions, and Content," *Canadian Journal of Philosophy* 43: 720-40.
- Bullot, Nicolas and Paul Égré (2010). "Objects and Sound Perception," *Review of Philosophy and Psychology* 1: 5-17.
- Burge, Tyler (2009). "Perceptual Objectivity," *The Philosophical Review* 118: 285-324.
- Byrne, Alex (2001). "Intentionalism Defended," *The Philosophical Review* 110: 199-240.
- Byrne, Alex and Heather Logue (2009). "Introduction," in *Disjunctivism: Contemporary Readings*, edited by Alex Byrne and Heather Logue, MIT Press.
- Calabi, Clotilde (2012). "Introduction," in *Perceptual Illusions: Philosophical and Psychological Essays*, edited by Clotilde Calabi, Palgrave Macmillan.
- Calvert, Gemma, Charles Spence, and Barry Stein (2004). *The Handbook of Multisensory Processes*, MIT Press.
- Campbell, John (2002). *Reference and Consciousness*, Oxford University Press.
- Chang, Edward, Jochem Rieger, Keith Johnson, Mitchel Berger, Nicholas Barbaro, and Robert Knight (2010). "Categorical Speech Representation in Human Superior Temporal Gyrus," *Nature Neuroscience* 13: 1428-32.
- Chisholm, Roderick (1957). *Perceiving: A Philosophical Study*, Cornell University Press.

Churchland, Patricia and Vilayanur Ramachandran (1994). "Filling In: Why Dennett is Wrong," in *Consciousness in Philosophy and Cognitive Neuroscience*, edited by Antti Revonsuo and Matti Kamppinen, Lawrence Erlbaum Associates.

Clark, Andy (2013a). "Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science," *Behavioral and Brain Sciences* 36: 181-253.

Clark, Andy (2013b). "Expecting the World: Perception, Prediction, and the Origins of Human Knowledge," *Journal of Philosophy* 110: 469-96.

Clark, Andy (2014). "Perceiving as Predicting," in *Perception and its Modalities*, edited by Dustin Stokes, Mohan Matthen, and Stephen Biggs, Oxford University Press.

Clark, Andy (2017). "Busting Out: Predictive Brains, Embodied Minds, and the Puzzle of Evidentiary Veil," *Noûs* 51: 727-53.

Coates, Paul (2007). *The Metaphysics of Perception*, Routledge.

Cohen, Michael, Daniel Dennett, and Nancy Kanwisher (2016). "What is the Bandwidth of Perceptual Experience?" *Trends in Cognitive Sciences* 20: 324-35.

Crane, Tim (2005). "What is the Problem of Perception?" *Synthesis Philosophica* 40: 237-64.

Crane, Tim (2006). "Is There a Perceptual Relation?" in *Perceptual Experience*, edited by Tamar Gendler and John Hawthorne, Oxford University Press.

Crane, Tim and Craig French (2015). "The Problem of Perception," *The Stanford Encyclopedia of Philosophy*.

de Gelder, Béatrice and Paul Bertelson (2003). "Multisensory Integration, Perception and Ecological Validity," *Trends in Cognitive Sciences* 7: 460-7.

Dennett, Daniel (1992). "Filling In Versus Finding Out: A Ubiquitous Confusion in Cognitive Science," in *Cognition: Conceptual and Methodological Issues*, edited by Herbert Pick, Paulus van den Broek, and David Knill, American Psychological Association.

Driver, Jon and Charles Spence (1998). "Attention and the Crossmodal Construction of Space," *Trends in Cognitive Sciences* 2: 254-62.

Driver, Jon and Charles Spence (2000). "Multisensory Perception: Beyond Modularity and Convergence," *Current Biology* 10: R731-5.

Driver, Jon and Charles Spence (2004). "Crossmodal Spatial Attention: Evidence from Human Performance," in *Crossmodal Space and Crossmodal Attention*, edited by Charles Spence and Jon Driver, Oxford University Press.

- Eagleman, David (2011). *Incognito: The Secret Lives of the Brain*, Random House.
- Fish, William (2004). "The Direct/Indirect Distinction in Contemporary Philosophy of Perception," *Essays in Philosophy* 5.
- Fish, William (2009). *Perception, Hallucination, and Illusion*, Oxford University Press.
- Fish, William (2013). "Perception, Hallucination, and Illusion: Reply to my Critics," *Philosophical Studies* 163: 57-66.
- Foster, John (2000). *The Nature of Perception*, Clarendon Press.
- French, Craig (ms). "The Problem of Illusion for Naïve Realism," unpublished manuscript. <http://philosophyofbrains.com/wp-content/uploads/2016/01/Craig-French-The-Problem-of-Illusion-for-Naive-Realism.pdf>
- French, Craig and Lee Walters (forthcoming). "The Invalidity of the Argument from Illusion," *American Philosophical Quarterly*.
- Friston, Karl (2010). "The Free-Energy Principle: A Unified Brain Theory?" *Nature Reviews Neuroscience* 11: 127-38.
- Frith, Chris (2007). *Making up the Mind: How the Brain Creates our Mental World*, Blackwell.
- Fujisaki, Waka, Shinsuke Shimojo, Makio Kashino, and Shin'ya Nishida (2004). "Recalibration of Audiovisual Simultaneity," *Nature Neuroscience* 7: 773-8.
- Geisler, Wilson (2008). "Visual Perception and the Statistical Properties of Natural Scenes," *Annual Review of Psychology* 59: 167-92.
- Genone, James (2014). "Appearance and Illusion," *Mind* 123: 339-76.
- Genone, James (2016). "Recent Work on Naïve Realism," *American Philosophical Quarterly* 53: 1-25.
- Genone, James (ms). "Comments on 'Naïve Realism and Illusion' by Boyd Millar," unpublished manuscript. <http://philosophyofbrains.com/wp-content/uploads/2016/01/James-Genone-Comments-on-'Naive-Realism-and-Illusion'-by-Boyd-Millar.pdf>
- Gladziejewski, Pawel (2016). "Predictive Coding and Representationalism," *Synthese* 193: 559-82.
- Grimes, John (1996). "On the Failure to Detect Changes in Scenes Across Saccades," in *Perception*, edited by Kathleen Akins, Oxford University Press.

Hairston, David, Donald Hodges, Jonathan Burdette, and Mark Wallace (2006). "Auditory Enhancement of Visual Temporal Order Judgment," *NeuroReport* 17: 791-5.

Harman, Gilbert (1990). "The Intrinsic Quality of Experience," *Philosophical Perspectives* 4: 31-52.

Hatfield, Gary (2007). "The Reality of Qualia," *Erkenntnis* 66: 133-68.

Hatfield, Gary (2016). "Perceiving as Having Subjectively Conditioned Appearances," *Philosophical Topics* 44: 149-78.

Herwig, Arvid and Werner Schneider (2014). "Predicting Object Features Across Saccades: Evidence from Object Recognition and Visual Search," *Journal of Experimental Psychology* 143: 1903-22.

Hickerson, Ryan (2004). "An Indirect Defense of Direct Realism," *The Journal of Mind and Behavior* 25: 1-5.

Hobson, Kenneth (2013). "In Defense of Relational Direct Realism," *European Journal of Philosophy* 21: 550-74.

Hohwy, Jakob (2007). "Functional Integration and the Mind," *Synthese* 159: 315-28.

Hohwy, Jakob (2013). *The Predictive Mind*, Oxford University Press.

Hohwy, Jakob (2016). "The Self-Evidencing Brain," *Noûs* 50: 259-85.

Hohwy, Jakob, Andreas Roepstorff, and Karl Friston (2008). "Predictive Coding Explains Binocular Rivalry: An Epistemological Review," *Cognition* 108: 687-701.

Howard, James, Jane Plailly, Marcus Grueschow, John-Dylan Haynes, and Jay Gottfried (2009). "Odor Quality Coding and Categorization in Human Posterior Piriform Cortex," *Nature Neuroscience* 12: 932-8.

Jackson, Frank (1977). *Perception: A Representative Theory*, Cambridge University Press.

Jousmäki, Veikko and Riitta Hari (1998). "Parchment-Skin Illusion: Sound-Biased Touch," *Current Biology* 8: R190.

Komatsu, Hidehiko (2006). "The Neural Mechanisms of Perceptual Filling-in," *Nature Reviews Neuroscience* 7: 220-31.

Kouider, Sid, Vincent de Gardelle, Jerome Sackur, and Emmanuel Dupoux (2010). "How Rich is Consciousness? The Partial Awareness Hypothesis," *Trends in Cognitive Sciences* 14: 301-7.

Kovács, Ilona, Thomas Papathomas, Ming Yang, and Ákos Fehér (1996). "When the Brain Changes its Mind: Interocular Grouping During Binocular Rivalry," *Proceedings of the National Academy of Sciences* 93: 15508-11.

Le Morvan, Pierre (2004). "Arguments Against Direct Realism and How to Counter Them," *American Philosophical Quarterly* 41: 221-34.

MacAdam, David (1985). *Color Measurement: Themes and Variations*, Springer-Verlag.

Macpherson, Fiona (2013). "The Philosophy and Psychology of Hallucination: An Introduction," in *Hallucination*, edited by Fiona Macpherson and Dimitris Platchias, MIT Press.

Macpherson, Fiona and Clare Batty (2016). "Redefining Illusion and Hallucination in Light of New Cases," *Philosophical Issues* 26: 263-96.

Margot, Christian (2009). "A Noseful of Objects," *Nature Neuroscience* 12: 813-4.

Martin, Michael (2001). "Beyond Dispute: Sense-Data, Intentionality and the Mind-Body Problem," in *History of the Mind-Body Problem*, edited by Tim Crane and Sarah Patterson, Routledge.

Martin, Michael (2002). "The Transparency of Experience," *Mind & Language* 17: 376-425.

Martin, Michael (2004). "The Limits of Self-Awareness," *Philosophical Studies* 120: 37-89.

Martin, Michael (2006). "On Being Alienated," in *Perceptual Experience*, edited by Tamar Gendler and John Hawthorne, Oxford University Press.

Martin, Michael (2010). "What's in a Look?" in *Perceiving the World*, edited by Bence Nanay, Oxford University Press.

Martin, Michael (2013). "Shibboleth: Some Comments on William Fish's *Perception, Hallucination & Illusion*," *Philosophical Studies* 163: 37-48.

Martin, Michael (2017). "Elusive Objects," *Topoi* 36: 247-71.

Masrouf, Farid (2017). "Space Perception, Visual Dissonance and the Fate of Standard Representationalism," *Noûs* 51: 565-93.

Matthen, Mohan (2016). "Is Perceptual Experience Normally Multimodal?" in *Current Controversies in Philosophy of Perception*, edited by Bence Nanay, Routledge.

Maund, Barry (2012). "Perceptual Constancies: Illusions and Veridicality," in *Perceptual Illusions: Philosophical and Psychological Essays*, edited by Clotilde Calabi, Palgrave Macmillan.

- McGurk, Harry and John MacDonald (1976). "Hearing Lips and Seeing Voices," *Nature* 264: 746-8.
- McLaughlin, Brian (2010). "The Representational vs. the Relational View of Visual Experience," in *The Metaphysics of Consciousness*, edited by Pierfrancesco Basile, Julian Kiverstein, and Pauline Phemister, Cambridge University Press.
- McLaughlin, Brian (2016). "The Skewed View from Here: Normal Geometrical Misperception," *Philosophical Topics* 44: 231-99.
- Metzinger, Thomas (2003). "Phenomenal Transparency and Cognitive Self-Reference," *Phenomenology and the Cognitive Sciences* 2: 353-93.
- Metzinger, Thomas (2009). *The Ego Tunnel: The Science of the Mind and the Myth of the Self*, Basic Books.
- Millar, Boyd (2015). "Naïve Realism and Illusion," *Ergo* 2: 607-25.
- Möttönen, Riikka, Christina Krause, Kaisa Tiippana, and Mikko Sams (2002). "Processing of Changes in Visual Speech in the Human Auditory Cortex," *Cognitive Brain Research* 13: 417-25.
- Murray, Micah and Christoph Herrmann (2013). "Illusory Contours: A Window onto the Neurophysiology of Constructing Perception," *Trends in Cognitive Sciences* 17: 471-81.
- Murray, Micah, David Lewkowicz, Amir Amedi, and Mark Wallace (2016). "Multisensory Processes: A Balancing Act Across the Lifespan," *Trends in Neurosciences* 39: 567-79.
- Nanay, Bence (2014). "Empirical Problems with Anti-Representationalism," in *Does Perception Have Content?* edited by Berit Brogaard, Oxford University Press.
- Nanay, Bence (2015). "The Representationalism versus Relationalism Debate: Explanatory Contextualism about Perception," *European Journal of Philosophy* 23: 321-36.
- Nanay, Bence (2016). "Philosophy of Perception: A Road Map with Lots of Bypass Roads," in *Current Controversies in Philosophy of Perception*, edited by Bence Nanay, Routledge.
- Navarra, Jordi, H. Henny Yueng, Janet Werker, and Salvador Soto-Faraco (2012). "Multisensory Interactions in Speech Perception," in *The New Handbook of Multisensory Processing*, edited by Barry Stein, MIT Press.
- Neale, Stephen (2017). "Silent Reference," in *Meanings and Other Things: Themes from the Work of Stephen Schiffer*, edited by Gary Ostertag, Oxford University Press.
- Nudds, Matthew (2009). "Recent Work in Perception: Naïve Realism and its Opponents," *Analysis* 69: 334-46.

O'Callaghan, Casey (2012). "Perception and Multimodality," in *The Oxford Handbook of Philosophy of Cognitive Science*, edited by Eric Margolis, Richard Samuels, and Stephen Stich, Oxford University Press.

O'Callaghan, Casey (2015). "Speech Perception," in *The Oxford Handbook of Philosophy of Perception*, edited by Mohan Matthen, Oxford University Press.

O'Callaghan, Casey (2016). "Enhancement Through Coordination," in *Current Controversies in Philosophy of Perception*, edited by Bence Nanay, Routledge.

O'Callaghan, Casey (2017). "Grades of Multisensory Awareness," *Mind & Language* 32: 155-81.

Odegaard, Brian, Min Yu Chang, Hakwan Lau, Sing-Hang Cheung (ms). "Inflation Versus Filling-in: Why We Feel We See More than We Actually Do in Peripheral Vision."
<https://www.biorxiv.org/content/early/2018/02/12/263244>

Olkkonen, Maria and Sarah Allred (2014). "Short-Term Memory Affects Color Perception in Context," *PLOS One* 9.

Olkkonen, Maria, Patrice McCarthy, and Sarah Allred (2014). The Central Tendency Bias in Color Perception: Effects of Internal and External Noise," *Journal of Vision* 14.

Olkkonen, Maria, Toni Saarela, and Sarah Allred (2016). "Perception-Memory Interactions Reveal a Computational Strategy for Perceptual Constancy," *Journal of Vision* 16.

Orlandi, Nico (2016). "Bayesian Perception is Ecological Perception," *Philosophical Topics* 44: 327-51.

Otten, Marte, Yair Pinto, Chris Paffen, Anil Seth, and Ryota Kanai (2017). "The Uniformity Illusion: Central Stimuli Can Determine Peripheral Perception," *Psychological Science* 28: 56-68.

Papineau, David (2014). "Sensory Experience and Representational Properties," *Proceedings of the Aristotelian Society* 114: 1-33.

Pautz, Adam (2013). "Do the Benefits of Naïve Realism Outweigh the Costs? Comments on Fish, *Perception, Hallucination and Illusion*," *Philosophical Studies* 163: 25-36.

Pautz, Adam (2016). "Experiences are Representations: An Empirical Argument," in *Current Controversies in Philosophy of Perception*, edited by Bence Nanay, Routledge.

Pessoa, Luiz, Evan Thompson, and Alva Noë (1998). "Finding Out About Filling-In: A Guide to Perceptual Completion for Visual Science and the Philosophy of Perception," *Behavioral and Brain Sciences* 21: 723-802.

Peters, Megan, Robert Kentridge, Ian Phillips, and Ned Block (2017). "Does Unconscious Perception Really Exist? Continuing the ASSC20 Debate," *Neuroscience of Consciousness* 1.

Petkov, Christopher, Kevin O'Connor, and Mitchell Sutter (2007). "Encoding of Illusory Continuity in Primary Auditory Cortex," *Neuron* 54: 153-65.

Phillips, Ian (2016). "Naïve Realism and the Science of (Some) Illusions," *Philosophical Topics* 44: 353-80.

Phillips, Ian (forthcoming). "Unconscious Perception Reconsidered," *Analytic Philosophy*.

Phillips, Ian and Ned Block (2016). "Debate on Unconscious Perception," in *Current Controversies in Philosophy of Perception*, edited by Bence Nanay, Routledge.

Preston, John and Severin Schroeder (2013). "The Neuroscientific Case for a Representative Theory of Perception," in *A Wittgensteinian Perspective on the Use of Conceptual Analysis in Psychology*, edited by Timothy Racine and Kathleen Slaney, Palgrave Macmillan.

Purves, Dale and Beau Lotto (2003). *Why We See What We Do: An Empirical Theory of Vision*, Sinauer.

Radonjic, Ana, Nicolas Cottaris, and David Brainard (2015). "Color Constancy in a Naturalistic, Goal-Directed Task," *Journal of Vision* 15.

Ramachandran, Vilayanur (1992). "Blind Spots," *Scientific American* 266: 86-91.

Robinson, Howard (1994). *Perception*, Routledge.

Robinson, Howard (2013). "The Failure of Disjunctivism to Deal with 'Philosophers' Hallucinations'," in *Hallucination*, edited by Fiona Macpherson and Dimitris Platchias, MIT Press.

Rogers, Brian (2014). "Delusions about Illusions," *Perception* 43: 840-5.

Rosenthal, David (1985). "Review of *Perception: A Representative Theory* by Frank Jackson," *Journal of Philosophy* 82: 28-41.

Russell, Bertrand (1912/1997). *The Problems of Philosophy*, Oxford University Press.

Schwartz, Robert (2016). "Perceptual Veridicality," *Philosophical Topics* 44: 381-403.

Sellars, Roy Wood (1939). "A Statement of Critical Realism," *Revue Internationale de Philosophie* 1: 472-98.

Sellars, Wilfrid (1956). "Empiricism and the Philosophy of Mind," in *Minnesota Studies in the Philosophy of Science, Vol. 1: The Foundations of Science and the Concepts and Psychology and Psychoanalysis*, edited by Herbert Feigl and Michael Scriven, University of Minnesota Press.

Shahin, Antoine, Christopher Bishop, and Lee Miller (2009). "Neural Mechanisms for Illusory Filling-In of Degraded Speech," *NeuroImage* 44: 1133-43.

Shams, Ladan, Yukiyasu Kamitani, and Shinsuke Shimojo (2000). "What You See is What You Hear," *Nature* 408: 788.

Shams, Ladan and Robyn Kim (2010). "Crossmodal Influences on Visual Perception," *Physics of Life Reviews* 7: 269-84.

Shimojo, Shinsuke, Yukiyasu Kamitani, and Shin'ya Nishida (2001). "Afterimage of Perceptually Filled-in Surface," *Science* 293: 1677-80.

Smart, John (2002). "The Compatibility of Direct Realism with the Scientific Account of Perception; Comment on Mark Crooks," *The Journal of Mind and Behavior* 23: 239-44.

Smith, David (2002). *The Problem of Perception*, Harvard University Press.

Smythies, John (2005). "How the Brain Decides what We See," *Journal of the Royal Society of Medicine* 98: 18-20.

Smythies, John (2009a). "Philosophy, Perception, and Neuroscience," *Perception* 38: 638-51.

Smythies, John (2009b). "'Reality' and 'Virtual Reality' Mechanisms in the Brain and their Significance," *Journal of Consciousness Studies* 16: 69-80.

Smythies, John (2011). "Reply to Bennett and Hacker," *Perception* 40: 252-4.

Smythies, John and Vilayanur Ramachandran (1997). "An Empirical Refutation of the Direct Realist Theory of Perception," *Inquiry* 40: 437-8.

Sollberger, Michael (2008). "Naïve Realism and the Problem of Causation," *Disputatio* 3: 1-19.

Sollberger, Michael (2012). "Causation in Perception: A Challenge to Naïve Realism," *Review of Philosophy and Psychology* 3: 581-95.

Sollberger, Michael (2015). "In Defence of a Structural Account of Indirect Realism," *European Journal of Philosophy* 23: 815-37.

Spence, Charles (2015). "Multisensory Flavor Perception," *Cell* 161: 24-35.

Sperber, Dan and Deirdre Wilson (2015). "Beyond Speaker's Meaning," *Croatian Journal of Philosophy* 15: 117-49.

Stein, Barry and Alex Meredith (1993). *The Merging of the Senses*, MIT Press.

Sumbly, William and Irwin Pollack (1954). "Visual Contribution to Speech Intelligibility in Noise," *Journal of the Acoustical Society of America* 26: 212-5.

Tong, Frank, Ken Nakayama, Thomas Vaughan, and Nancy Kanwisher (1998). "Binocular Rivalry and Visual Awareness in Human Extrastriate Cortex," *Neuron* 21: 753-9.

Tye, Michael (1992). "Visual Qualia and Visual Content," in *The Contents of Experience*, edited by Tim Crane, Cambridge University Press.

Tye, Michael (1995). *Ten Problems of Consciousness*, MIT Press.

Watkins, Susanne, Ladan Shams, Sachiyo Tanaka, John-Dylan Haynes, and Geraint Rees (2006). "Sound Alters Activity in Human V1 in Association with Illusory Visual Perception," *NeuroImage* 31: 1247-56.

Watkins, Susanne, Ladan Shams, Oliver Josephs, and Geraint Rees (2007). "Activity in Human V1 Follows Multisensory Perception," *NeuroImage* 37: 572-8.

Weil, Rimona and Geraint Rees (2011). "A New Taxonomy for Perceptual Filling-in," *Brain Research Reviews* 67: 40-55.

Westerhoff, Jan (2016). "What it Means to Live in a Virtual World Generated by Our Brain," *Erkenntnis* 81: 507-28.

Wilcox, Stephen and Stuart Katz (1984). "Can Indirect Realism be Demonstrated in the Psychological Laboratory?" *Philosophy of the Social Sciences* 14: 149-57.

Wokke, Martijn, Annelinde Vandenbroucke, Steven Scholte, and Victor Lamme (2013). "Confuse Your Illusion: Feedback to Early Visual Cortex Contributes to Perceptual Completion," *Psychological Science* 24: 63-71.

Wright, Edmond (1986). "Wilcox and Katz on Indirect Realism," *Philosophy of the Social Sciences* 16: 107-13.

Yarrow, Kielan, Patrick Haggard, Ron Heal, Peter Brown, and John Rothwell (2001). "Illusory Perceptions of Space and Time Preserve Cross-Saccadic Perceptual Continuity," *Nature* 414: 302-5.

Young, Benjamin (2016). "Smelling Matter," *Philosophical Psychology* 29: 520-34.

Zieba, Pawel (2017). "Naïve Realism about Unconscious Perception," *Synthese*.