The Representational Theory of Consciousness

David Bourget

A thesis submitted for the degree of

Doctor of Philosophy

of the Australian National University

January 2010

Acknowledgments

I've had the best panel ever. David Chalmers, Chair of my panel, has been immensely generous with his time. His prompt, extensive, regular feedback on my dissertation and ancillary projects has been transforming. Daniel Stoljar has also been extremely helpful. Each opportunity I've had to discuss my work with Frank Jackson has been highly profitable.

My fellow students have also helped me. I've learned a great deal through the "R2D2" meetings with Michael Lachelt, Alex Lamey, Tom Rand, Brad Richards, and Joel Walmsley at UofT. At ANU, I've benefited from the help of Ben Blumson, Magdalena Balcerak Jackson, Farid Masrour, and Dan Marshall at critical junctures.

A good number of ANU visitors and staff members outside my panel have contributed to this project with their advice, comments, criticisms, or well-timed incredulous stares, often in ways they might not have realized. In alphabetical order: Tim Bayne, Berit Brogaard, Alex Byrne, Tim Crane, Kati Farkas, William Fish, Benj Hellie, Christopher Hill, Terry Horgan, Uriah Kriegel, Stephan Leuenberger, Fiona Macpherson, Adam Pautz, Joe Salerno, Susanna Schellenberg, Wolfgang Schwarz, Susanna Siegel, Nico Silins, Declan Smithies, Paul Snowdon, Brad Thompson.

Finally, thanks to William Seager, Ronald de Sousa, and Gurpreet Rattan for the guidance they provided me as I was beginning to articulate this project at UofT.

Abstract

The representational theory of consciousness (also known as "representationalism") could potentially supply a framework in which to formalize the individual natures of conscious states, their inter-relations, and the relations they stand in to physical states. But this theory is itself in need of severe regimentation and a rigorous defense before it can begin to play this role as part of scientific inquiry. This is what I aim to achieve with this thesis.

I begin by addressing certain difficulties with the notion of representational content. I then isolate the core tenets of representationalism by subtracting from it all commitments which are not essential to its most important applications, including in particular its potential use as a framework for the scientific study of consciousness. The result is a thesis I call "virtualism".

My case for virtualism proceeds in three steps. I first argue that states of sensory consciousness can be described using perceptual verbs intensionally, as when one says that one "sees stars". I then use perceptual verbs in this way to show that a restricted form of virtualism is true of sensory consciousness. Finally, I argue from this restricted virtualism to the full virtualist theory on the basis of the phenomenological unity of consciousness and other considerations.

The last part of the thesis addresses central objections to representationalism. The objections I consider fall into two broad categories: those that target the representationalist claim that the phenomenology of consciousness is exhausted by its representational contents, and those that rest on the view known as "naive realism".

Contents

Introduction

A satisfactory solution to the problem of consciousness would take the form of a simple yet fully general model that specifies the precise conditions under which any given state of consciousness occurs. Science has uncovered numerous correlations between consciousness and neural activity, but it has not yet come anywhere close to this. We are still looking for the Newtonian laws of consciousness.

One of the main difficulties with consciousness is that we lack a language in which to formulate illuminating generalizations about it. Philosophers and scientists talk about "what it's like", sensations, feelings, and perceptual states such as seeing and hearing. This language does not allow a precise articulation of the internal structures of conscious states and their inter-relations. It is inadequate to capture relations of the kind we are looking for between conscious states and physical states.

In this thesis I refine and defend a theory of consciousness which promises to solve this regimentation problem: the representational theory of consciousness. I argue that the representational theory can solve the regimentation problem and smooth out other important obstacles to a fruitful study of consciousness. I also make a case for the theory independently of its payoffs, and I discuss the leading opposing theories at some length.

In the rest of this introduction, I will clarify what I mean by "consciousness", provide an initial characterization of the representational theory, and outline my project in more detail.

Consciousness

At this stage in the thesis, I cannot provide a precise and unambiguous definition of consciousness which will please everyone, but I will do my best to fix ideas. What matters is that you and I attend to the same phenomenon. Articulating the nature of this phenomenon is the project the rest of this thesis tackles.

The kind of consciousness this thesis is about is the kind which is widely thought to pose a unique challenge to scientific explanation. It is the kind which has come to be called "phenomenal consciousness". To be phenomenally conscious is to instantiate a *phenomenal state*. The paradigmatic phenomenal states are states we instantiate in the course of our sensory, emotional, and cognitive experiences. Here I am using "experience" in a sense which should be familiar from everyday talk. I will give some examples of each kind of experience just mentioned.

At the moment I have a mild back pain. My experience of pain is a paradigmatic sensory experience. There are many other kinds of sensory experience. When I look around myself, for example, I undergo a large number of visual experiences. Visual experiences, as I think of them, have a felt component a little bit like pain experiences. As Nagel (1974) puts it, there is something it's like for a subject to have a visual experience. There are also auditory, tactile, olfactory, gustatory, kinaesthetic, and proprioceptive experiences, and probably many other kinds of sensory experience which do not fall in any of these categories. There are probably infinitely many kinds of possible sensory experience possible creatures could have. All have felt components which are part of their essences as sensory experiences.

Emotional experiences also come in a wide variety. For example, one undergoes emotional experiences when one feels anxious, sad, relieved, or elated. Emotional experiences should not be conflated with emotions. On one common understanding of emotions, at least, they are states which can persist independently of how one feels. For example, one can be angry at a time without actually feeling angry at that time. We can leave the question of how emotions and emotional experiences

relate to each other largely open for now, but we need to recognize that emotional experiences, with their felt components, do not always accompany the emotions with which they tend to be associated. It is emotional experiences which are of interest here, not emotions.

I count as cognitive all conscious experiences which are not normally associated with emotions or sensory processes. Cognitive experiences are more elusive than sensory and emotional experiences. Goldman (1993) draws attention to cognitive experiences by comparing what it is like to hear and understand a statement with what it is like to hear the same statement without understanding it. There is a certain feeling of understanding missing in the second case—a feeling Strawson (1994) describes as an *understanding experience*. Goldman also mentions the tip-of-the-tongue feeling. There are many other kinds of "cognitive feeling", for example, the feeling that something is right (or not right), the feeling of being confused, and the feeling of deja vu. We constantly rely on cognitive feelings like these for guidance in everyday life.

Also central to one's mental life is sensory imagination (which I count as cognitive, but that is merely a convenient terminological choice). All of us have some capacity for visual sensory imagination. For instance, one can faintly visualize (experience) an object one is looking at being moved from its actual location to another location. Aural imagery is also very common. When you talk to yourself "in your head", you are experiencing aural imagery. I will leave other, more controversial forms of sensory imagery aside for now.

There might be cognitive experiences which are less sensation-like than emotional feelings and sensory imagery. We could describe these as *pure phenomenal thoughts*. I will remain neutral as to whether there are pure phenomenal thoughts.

Each kind of experience I have mentioned has a felt component: for each kind of experience I have mentioned, there is something it is like to have an experience of this kind. Phenomenal states are states of the kind one is in when undergoing

sensory, emotional, and cognitive experiences, in virtue of which there is something it is like to have these experiences.

Phenomenal state A state of the kind best exemplified by the states a) instantiated by individuals in sensory, emotional, and cognitive experiences b) individuated by the felt components they confer to such experiences.

It is noteworthy that experiences, in the everyday sense of "experience", are not the same as phenomenal states or instantiations of phenomenal states. Consider the definition of the verb "to experience" given by the Oxford English Dictionary. According to the OED, to experience something is to "encounter" or "undergo" an "event or occurrence".[1] It is clear that this is how "to experience" is used when one says (for example) that Microsoft has experienced a slowdown. It is used in exactly the same way when talking about consciousness-involving sensory episodes in a lay context.[2] Take for example these three everyday statements:

(1) I experienced pain

(2) I experienced flu symptoms

(3) I experienced powerlessness

The OED interpretation of "to experience" makes sense of all three: I encountered / underwent pain, I encountered / underwent flu symptoms, and I encountered / underwent powerlessness.

On this everyday interpretation of "experience", a sensory experience of red is a kind of encounter with redness. A sensory experience of an object (say, my kitchen table) is a kind of encounter with that object. Qua encounter, an experience requires the presence of what is encountered: one could not possibly encounter my kitchen table without being related to it in some way. The same goes for an experience of

[1] "experience noun", The Oxford Dictionary of English (revised edition). Oxford University Press, 2005. Oxford Reference Online. Retrieved on 2 May 2009 from: http://www.oxfordreference.com/views/ENTRY.html?subview=Main&entry=t140.e26250

[2] Thanks to Alex Byrne for helpful comments here.

Sirius (the star): one could not possibly have a visual experience of Sirius in the OED's sense without being affected by Sirius in some way.

Many theorists would deny that there is such a thing as a Sirius-feeling. There is a phenomenal state one typically instantiates when one visually experiences Sirius in the OED's sense (when one visually encounters Sirius), but this state is not characteristic of Sirius experiences (as opposed to experiences of other stars). According to the theorists in question, one could instantiate exactly the same phenomenal state as part of a visual experience of (a visual encounter with) another star. For this reason, we cannot say simply that the phenomenal states are the states instantiated in sensory, emotional, and cognitive experiences. This would invite an understanding of phenomenal states as encounter states, and such states clearly involve more than what the philosophers in question take phenomenal states to be by definition, namely, states individuated by their felt components (their "what it's like" aspect). This is why I stipulate that phenomenal states are individuated by their felt components: because the felt aspect of an experience in the OED sense need not involve the particular object of the experience.

We will see later in the thesis that there are some who maintain (contrary to the aforementioned philosophers) that many if not all phenomenal states essentially have external particulars as components. This claim would be nearly trivial if phenomenal states were the states which constitute experiences in the everyday sense. On the other hand, it is far from trivial on my definition of phenomenal states. But my definition is not meant to exclude this view. The aim is to isolate the substantive issue which is at stake in the debate between proponents of this view and their opponents.

So far I have been using the term "experience" in its everyday sense. It is not used in this way by everyone, and the everyday sense is rather nebulous (though clear enough for the use to which I put it, I hope). For our purposes, it is helpful to restrict and regiment our use of "experience". From now on, I will use this term as

follows:

Experience An event which consists in instantiating a phenomenal state.

This definition of experiences makes them events of the kind described by Kim's (1976; 1991) theory of events. I think this is a common understanding of the noun "experience" among philosophers, but there are probably other uses. We must keep this in mind as we progress. We must also precisify the meaning of the verb "to experience". As I use it, it means instantiating a phenomenal state.

Theorists often talk about "phenomenal characters" and "phenomenal properties". Here I am going to use these terms interchangeably with "phenomenal state" (I take states to be properties). I take it that my use of these terms is common, but there are no doubt exceptions.

The representational theory

The representational theory of consciousness also goes by the names of "representationalism" and "intentionalism". The best known versions of this theory are illustrated by the following statements:

- (1) All mental facts are representational facts, and (2) all representational facts are facts about informational functions. (Dretske 1995: xiii)
- Phenomenal character is one and the same as representational content that meets certain further conditions. (Tye 2000: 45)
- The propositional content of perceptual experiences in a particular modality (for example, vision) determines their phenomenal character. (Byrne 2001)
- Phenomenal properties are identical to certain representational properties. (Chalmers 2004)

· Qualia are actually intentional contents, represented properties of represented objects. (Lycan 2005)

To a first approximation, the points which unify the representational theories of the preceding authors are a) that experiences essentially have representational or intentional contents and b) that their phenomenal characters are largely determined by their contents.

The recent history of the representational theory can be traced back at least to Anscombe (1965), Armstrong (1968), Hintikka (1969), and Pitcher (1971). It has attracted a significant following during the past decades, but it remains controversial.[3] A large tangle of objections, counterarguments, counter-counterarguments, misunderstandings, and alternate formulations has developed. Unless some systematic way of simplifying and clarifying the debate is discovered, there is little hope for further convergence. This is part of what I aim to provide in this thesis.

It strikes me that the bulk of the objections to representationalism are only effective against specific versions of the theory. Some oppose the externalist commitments of Dretske's, Lycan's, and Tye's views, but there are internalist versions of representationalism. Others object that experiences do not have satisfaction conditions, but we don't have to think of the representational contents of experiences as satisfaction conditions. Some argue that experiences cannot be individuated by how the world "looks" or "seems" to one in them, but that experiences make the world look or seem a certain way in a non-trivial sense is not a general commitment of

[3] According to the PhilPapers Survey 2009 (http://philpapers.org/surveys/), the distribution of relevant views among the faculties of leading philosophy departments is as follows:

· representationalism: 31.4%

· qualia theory: 12.2%

· disjunctivism: 10.9%

· sense-datum theory: 3.1%

· other theories (e.g. the theory of appearing): 3.8%.

While representationalism is the most widely accepted view, the debate is far from settled. It is also noteworthy that the theory defended here is stronger than the average representational theory in many respects (though weaker in other respects).

representationalism. And on it goes.[4]

Since there seems to be no common basis for these objections, one naturally wonders whether we would be left with anything of interest after subtracting everything that is controversial from representationalism. Maybe not, but I believe that it is possible to articulate a theory which captures the essence of the representationalist program yet is free from the most controversial implications of current versions of the theory. Relatedly, I believe that it is both possible and necessary to clarify this program. One of the main reasons critics pick on the details of particular representationalist views is that there is no widely agreed upon statement of the core theory representationalists share. In particular, there is a bewildering variety of interpretations of the notion of representational content apparently central to the theory. As it is, one could reasonably doubt that there is such a thing as *the* representational theory of consciousness. I believe that there is a core representational theory, and I believe that this theory has few of the features of specific representational theories which have been found controversial. My goal with this thesis is to articulate this theory, exhibit its potential as a framework for the scientific study of consciousness, build a positive case for it, and show how well it fares on key issues which have surfaced in the debates surrounding representationalism.

I call the purified representationalism I defend *virtualism*. I use this new term for two reasons. Firstly, the names "representationalism" and "intentionalism" have by now received so many divergent interpretations that their linguistic meanings could well be beyond repair. Secondly, whether virtualism really is a kind of representationalism or not does not matter; I don't want to invite a debate on this by using the term. What matters (to me, anyway) is that virtualism captures what representational theories share with one another, including in particular their applications.

It would help to have in sight some of these applications before starting. I already mentioned representationalism's promise as a solution to the problem of regi-

[4] See section 1.2 for more on these objections.

menting consciousness. In my opinion, this is the most important application of the theory. Let suppose for the purposes of illustration that a phenomenal state is a state in which one phenomenally represents a certain proposition. Let us not worry about what it is to phenomenally represent a proposition; this will become clear later. If a representational thesis of this type were correct, all differences between phenomenal states would be accounted for by differences in the propositions one represents in them (their "contents"). Now, propositions can be described using logic, and general relations between propositions and others entities, be they other propositions or worldly entities, can likewise be drawn formally. For this reason, it is reasonable to expect this kind of representationalism to enable a new level of regimentation in describing phenomenal states, their inter-relations, and their place in nature. I will elaborate more on this in chapter 2.

Representationalism also promises to help solve (or dissolve) the problem of perception. The problem of perception is best seen as a tension between two observations. On the one hand, perceptual experience appears to afford us a direct, unmediated awareness of our surroundings. On the other hand, no one can deny the possibility of illusions and hallucinations. These two observations come in tension through arguments sense-datum theorists have offered which purport to show that the possibility of illusions and hallucinations implies that perceptual awareness is primarily awareness of mental "sense data" (c.f. Ayer 1940; Moore 1905; Jackson 1976,1977). If sense-datum theorists were correct, the two aforementioned observations would be inconsistent. For many philosophers of perception today, the problem of perception is to show that sense-datum theorists are wrong (Crane 2005). To show this, we need to make room for both the immediacy of perceptual experience and the possibility of illusion and hallucination. The sense-datum theory is sometimes referred to as "the representative theory of perception", but it should not be conflated with representationalism, because the latter is generally intended to provide an alternative to the sense-datum theory which achieves just this.

How does representationalism undermine the inference from the possibility of illusions and hallucinations to the conclusion that what we are immediately aware of in perceptual experience are mental sense data? Representationalists hold that experience is a species of representation. This does not mean that our experience of the world is "mediated" in the sense that an experience of the world can only be had by experiencing some internal entity which stands for or represents external facts. To a first approximation, all this means is that one can experience an F without there being an F. This is on the face of it perfectly compatible with our having experiences of everyday objects and qualities without those experiences being mediated by awareness of anything like sense data. Compare with the case of belief: one can believe that there is an F without there being an F, but this does not imply that belief is mediated by awareness of anything like sense data. As we will see in chapter 2, the argument from illusion relies on an assumption to the effect that one cannot experience an F unless one is suitably related to an F. Since representationalists' claim that experience is a species of representation is incompatible with this assumption, the representationalist view blocks the argument from illusion. Moreover, it does this without rejecting the immediacy of experience. It therefore seems to dissolve the apparent dilemma posed by the case for sense data.

In addition to the sense-datum theory, there are two main incompatible alternatives to representationalism today: the intrinsic qualia theory and disjunctivism.[5] According to the former, consciousness is a matter of a state or person instantiating intrinsic qualitative properties of a special kind. The principal motivations for the intrinsic qualia theory will be discussed in chapter 6. To a first approximation, disjunctivists claim that there are two kinds of phenomenal state: those which can occur in veridical conditions only, and those which can occur in other conditions. I will discuss disjunctivism at some length in chapters 7 and 8.

[5]Of course, these are by no means the only alternatives to representationalism. For example, there is also the "the theory of appearing" (Langsam 1997; Alston 1999). See footnote 3 for the distribution of views in philosophy of perception.

Outline

Part I. The theory and its applications

Chapter 1 articulates the kind of representationalism I aim to defend (virtualism). I begin by highlighting the difficulties which led me to concentrate on the virtualist view. The alternative representational theories all face one of two problems: either they are too vague to yield a good framework for the scientific study of consciousness or they are loaded with questionable commitments. Virtualism minimizes both difficulties, or so I try to show.

Chapter 2 argues that virtualism can play the key explanatory roles widely attributed to representationalism. This chapter also aims to provide an initial motivation for the theory by making these roles explicit.

Part II. A case for virtualism

Chapter 3 puts forward a framework I refer to as *the perceptual conception of sensory experience.* The perceptual conception of sensory experience enables us to describe sensory phenomenal states precisely in everyday language by using perceptual verbs intensionally. This will be important in the chapters which follow.

Chapter 4 builds on the perceptual conception of sensory experience to make a case for a restricted version of virtualism I call *sensory virtualism*. I offer an account of the semantics of perceptual verbs in the process.

Chapter 5 argues that we should not stop at sensory virtualism: if sensory virtualism seems plausible, so should the full virtualist theory. This chapter includes a discussion of objections to representationalism which turn on the role of sensory modalities.

Part III. Alternatives

Chapter 6 considers objections to representationalism which tend to support the case for a qualia theory. I discuss the objections from perspective, blurry vision, double vision, imagery, inverted spectra, and other cases or phenomena which have been held to pose a challenge for representationalism and require that we postulate intrinsic qualia. These matters have been extensively discussed in the literature already; I limit myself to novel responses the framework developed here enables.

Chapter 7 assesses the phenomenological evidence for and against disjunctivism. Disjunctivists have often claimed that disjunctivism is supported by the "naive" conception of experience we acquire through introspection. In this chapter, I argue that virtualism, not disjunctivism, is supported by introspection.

Chapter 8 complements the preceding chapter with a discussion of two central motivations for disjunctivism which are largely independent of the revelations of introspection. I first ask whether disjunctivism is justified by considerations pertaining to skepticism and the role of experience in grounding thoughts about the external world. I then argue that disjunctivism should be rejected on the ground that perceptual experience is massively illusory.

Part I

The theory and its applications

Chapter 1

Core representationalism

In introduction I glossed representationalism as the view that a) experiences essentially have representational or intentional contents and b) their phenomenal characters are largely determined by their contents. I said that I am looking for the core representational theory—the ecumenical representational theory which lends itself best to the applications representationalism is generally supposed capable of. The foremost application I am concerned with is to provide a framework for the scientific study of consciousness. The first thing I want to do in this chapter is to show that the conjunction of (a) and (b) is not suitable. This statement of representationalism is unsatisfactory because the generic notion of intentionality or representation it relies upon is not amenable to sufficient clarity and precision. I will also consider the alternatives which have been offered to statements along the lines of (a) and (b). I will end the chapter by introducing the formulation of core representationalism I favor, the theory I call *virtualism*.[1]

[1] A good part of the first section of this chapter comes from my article "Consciousness is underived intentionality", forthcoming in *Noûs*.

1.1 Intentionality, attitudes, and representation

The statement of representationalism I have been using so far is satisfactorily ecumenical and non-committal. But what are the intentional contents referred to in this statement?

> Every mental phenomenon is characterized by what the Scholastics of the Middle Ages called the intentional (or mental) in-existence of an object, and what we might call, though not wholly unambiguously, reference to a content, direction toward an object (which is not to be understood here as meaning a thing), or immanent objectivity. Every mental phenomenon includes something as object within itself, although they do not all do so in the same way. (Brentano 1874: 88-89)

Brentano's articulation of the concept of intentionality is a landmark in the history of philosophy and psychology, but it is insufficiently clear for our purposes. We are looking for a representational theory that could serve as a conceptual foundation to formulate precise, clear, and general hypotheses about consciousness. We cannot let our theory rest on a metaphor such as "direction toward an object". It will also not do for our purposes to explain intentional states in terms of aboutness, or to say that they have aboutness and other features (aspectuality, intensionality, etc). Aboutness is just that pointing beyond or directedness Brentano is gesturing toward.

Intentional states are sometimes equated with propositional attitudes. So an alternative approach to representationalism is to say that phenomenal states are propositional attitudes. But "propositional attitude" is also a term of art whose use is not tied to any clear definition. I suspect that there would be much disagreement as to whether the following are propositional attitudes: dreaming that P, knowing that P, being frustrated that P, relishing the fact that P, remembering that P, forgetting that P, neglecting that P, being blind to the fact that P, being unable to believe that P, seeing that P, mistakenly thinking that P, consciously wishing that P, having heard

that P, deceiving oneself about the fact that P, supposing that P, routinely proving that P, never saying that P, being undecided about the claim that P, being born before anyone said that P. There are also states which are not naturally attributed using "that"-clauses but might be thought to fall in the category of propositional attitudes, e.g. looking for help, searching for the exit, trying to speak, liking proposition #2, believing many propositions, etc. The fuzzy set which is the set of propositional attitudes has such a large fringe, it is hard to give determinate content to the claim that phenomenal states are propositional attitudes.

There is of course a set of states which everyone would agree are propositional attitudes. I would say that these *canonical propositional attitudes* include belief, desire, and states which appear to be intensity variations on or complex combinations of these states, for example, conviction, suspicion, wish, hope, etc. One could use canonical propositional attitudes in either of two ways to clarify the notion of a propositional attitude. First, one could stipulate that the propositional attitudes are the canonical propositional attitudes. Alternatively, one could define propositional attitudes as states which are in some sense like the canonical propositional attitudes. Either way, one could then equate the claim that phenomenal states are intentional with the claim that they are propositional attitudes.

The problem with the first approach is that it seems fairly clear that phenomenal states are not canonical propositional attitudes. Imagine for example that you are looking at a stick which is half immersed in water. Even if your experience in some sense presented the stick as bent to you, you would not necessarily believe, wish, suppose, desire, or intend it to be bent, or have any other canonical propositional attitude toward its being bent. Alternatively, suppose that you are lying in bed with your eyes closed and an image of a beach pops in your mind. It does not follow from the fact that you have an experience of a beach that you have any canonical propositional attitude toward any proposition about any beach.

Arguably the best potential examples of phenomenal states which are also propositional attitudes are felt urges. Imagine for example that Bob felt a strong urge to eat a piece of black forest cake someone put in front of him. This feeling would be a phenomenal state. It might also seem to be a kind of propositional attitude—a desire-like propositional attitude. Similarly, one might think that a perceptual experience accompanied by a certain feeling of conviction ought to count as a belief-like propositional attitude.

But imagine that Bob felt an urge to eat a piece of cake while having no corresponding cognitive or behavioral dispositions: although he felt an urge to eat it, he gave strong nonverbal signs of loathing it, asked the person next to him if she would not want his share, and was profoundly relieved he would not have to suffer its sweetness when she accepted it. The natural thing to say in this case would be that Bob momentarily felt as if he wanted to eat the cake but did not really want to. The same is true of belief-like propositional attitudes and the feelings which sometimes accompany them. For instance, a claim can momentarily feel right to you without your believing it. One can feel as if one believed something without really believing it. Although feelings are normal concomitants of propositional attitudes, they don't seem to be canonical propositional attitudes.

If phenomenal states are not canonical propositional attitudes, a representationalist can at best say that they are *like* canonical propositional attitudes. This takes us back to the notion of intentionality, because the alternative to defining intentionality through the aboutness metaphor is to give examples of it, and one would have to use the canonical propositional attitudes as examples.[2] So at this stage the

[2]It is noteworthy that Brentano uses experiences as his primary examples of intentional states (of mental states, in fact, but he means to equate the two). Somehow, propositional attitudes have become the canonical intentional states. Brentano's true explication of the notion of intentionality is a list of examples, and the main ones are "presentations". Brentano explains the notion of a presentation as follows:

> By presentation I do not mean that which is presented, but rather the act of presentation. Thus, hearing a sound, seeing a colored object, feeling warmth or cold, as well as similar states of imagination are examples of what I mean by this term. (Brentano 1874: 479)

question is whether the notion of intentionality or propositional attitude which the representationalist wants to apply to experience could be clarified sufficiently by using the canonical propositional attitudes as examples. We are now getting closer to how the representational theory is normally understood in practice. Even though some explicit definitions of "intentionality" and "propositional attitude" have been attempted, it seems clear that the terms are in practice grounded in the examples of canonical propositional attitudes.

There seems to be three possible ways of defining intentionality (or propositional attitudes) by pointing to the canonical propositional attitudes as examples:

> MOST: Intentionality is the most specific kind which encompasses all canonical propositional attitudes.
>
> LEAST: Intentionality is the least specific kind which encompasses all canonical propositional attitudes.
>
> NATURAL: Intentionality is the most specific, interesting natural kind which encompasses all canonical propositional attitudes.

LEAST is clearly a nonstarter: it picks out the vacuous Entity or State kind.

The problem with MOST is that many properties a representationalist should probably not ascribe to phenomenal states happen to be had by all canonical propositional attitudes. Here are some examples:

(1.1) Not being a phenomenal state

(1.2) Being at least partially a functional or dispositional state

(1.3) Being "cognitive" and "conceptual" (on some understanding of these terms)

(1.4) Having a content that can in principle be expressed in plain English

---

It would not be completely out of line with the long-term tradition to stipulate that experiences are canonical intentional states. The only thing is that this would trivialize much of the representationalist position. I am interested in the non-trivial theory which appears to have interesting applications, so I will set this option aside.

I have already argued that phenomenal states are not canonical propositional attitudes, and from this it follows that canonical propositional attitudes are not phenomenal states (1.1). I have not said anything about 1.2, however, and the point will be important later, so I will say something about it now. I will leave the claim that all canonical propositional attitudes have properties 1.3 and 1.4 unargued; two good examples should suffice.

That the canonical propositional attitudes are essentially functional or dispositional (property 1.2) is easily seen in light of the fact that experiences are not canonical propositional attitudes. Take the phenomenal states I used as examples of phenomenal states which are not canonical propositional attitudes. Assuming for now that they have content in some sense or other—they certainly have some kind of directedness, in some vague sense of this—, what is missing from them to confer propositional attitudes? What else would it take for you to count as believing that the stick is bent, or as desiring to be at the beach? Dispositions to behave (move) in certain ways would seem insufficient. For not believing that P seems entirely consistent with experiencing P and having exactly the behavioral dispositions of someone who believes that P. It is a familiar point that the very same behavioral dispositions can be explained just as well by mutually exclusive sets of attitudes, e.g., a set involving the belief that P and other propositional attitudes, and a set involving the belief that not-P among other propositional attitudes. But if behavioral dispositions would not be sufficient for you to believe that the stick is bent or desire being at the beach, what is missing has to be dispositions to token other mental states. Perhaps you would have to be robustly disposed to feel certain things, e.g. a certain conviction for belief, a certain pull for desire. Or perhaps you would need dispositions to token propositional attitudes consistent with believing that the stick is bent or desiring to be at the beach. Or one might think that either robust dispositions to have relevant feelings or dispositions to token relevant propositional attitudes would do it. Whatever the case may be, these examples bring out the fact that believing

and desiring are attitudes *in the ordinary sense*—in the same sense in which being annoying and being forward are attitudes. Your beliefs and desires are not just static representations in your brain. They also involve (if they are not exhausted by) complex dispositions to token certain patterns of mental states (phenomenal states and other propositional attitudes) and (perhaps) certain patterns of behavior. As such, they are *in part* functional or dispositional.[3] I will not argue more for this point, because I take it to be as close to orthodoxy as there ever is in the field.

Arguably, none of properties 1.1-1.4 are essential to phenomenal states.[4] But if, as seems plausible, all the canonical propositional attitudes have these properties, then plausibly all states of the most specific kind which encompasses all the canonical propositional attitudes have these properties, too. So option MOST does not seem to yield a notion of intentionality that is suitable for a representationalist's purposes, though it could of course be suitable for other purposes.[5]

The main problem with NATURAL is that there seems to be many equally interesting and specific natural kinds which encompass the canonical intentional states specified. For example, there is the kind Millikan (1984; 1989; 2004) identifies with representation—roughly, states which must correspond in certain ways to external facts for the systems which "consume" them to perform normally in a biological sense. There is also the kind Dretske (1995; 2003) identifies with representations, that is, states which have the function of indicating external facts. And there is the kind of intentionality which is constituted by interpretability—the sort described by Davidson's (1973; 1974; 1984) and Dennett's (1971; 1987) views. Any "theory

[3]Note how weak this claim is. Since it only implies that certain relations to distinct mental states are partly constitutive of attitudes, it is not open to objections to functionalism which purport to show that functional role is insufficient for attitudes. For example, it is not open to Searle's (1980) Chinese room and Block's (1978) China brain objections. It is also not open to Galen Strawson's (1994) Weather Watcher argument against "neobehaviorism", because it does not draw any necessary connection between mental states and behavior.

[4]See Bourget (forthcoming) regarding 1.2.

[5]Some objections to representationalism seem to arise from the assumption that something like MOST defines intentional states. Robinson's (1994) claim that experience is not a kind of representation because representations are "cognitive" and don't have "feel" seems to be an example of this.

of intentionality" which predicts that there is intentionality where there are canonical propositional attitudes potentially describes a highly specific natural kind that encompasses all these attitudes, and all seem equally striking or interesting. NATURAL therefore leaves intentionality talk highly indeterminate. It also retains the problem we had with MOST. For example, all the canonical propositional attitudes appear to be non-phenomenal states, so it might well be that the most specific, interesting natural kind which encompasses all of them (if any) is a kind of state whose instances are necessarily non-phenomenal.

The preceding problems bring out the fact that we cannot escape specifying explicitly what it is that makes a state intentional. A definition by ostension works when vagueness is acceptable or some abstraction method is specified implicitly or explicitly. Often, the abstraction method is to focus on the most striking or interesting kind which encompasses the canonical examples (that is roughly how Putnam 1975 claims we pick out natural kinds). But in the case of intentionality there seems to be no abstraction method available that would yield a concept usable as part of the representationalist framework.

Compared to "intentionality" and "propositional attitude", the term "representation" has the advantage of grounding in everyday language. Examples of representation in the everyday sense include diagrams, street signs, photographs, semaphore signs, sculptures, equations, computer models, and sentences. Contrary to "intentionality", there seems to be a pre-theoretic kind this term latches on.

The problem with saying that experiences are representations in the everyday sense is that this kind appears to have a feature experiences clearly don't have: it appears to be a kind of thing whose instances are dependent on the intentions of individuals for their existence. A colored sheet of glossy paper, for example, cannot be a photograph unless it has been produced by an apparel intended to photograph or duplicate photographs. The sheet's existence does not depend on the intentions of individuals, but its being a photograph or, put differently, the existence of a pho-

tograph, depends on the intentions of individuals. Likewise, a string of marks on a sheet of paper is not a sentence (i.e. a meaningful entity) unless it or some kind it exemplifies is intended to be by an individual or a collective, and a white arrow on a black background is not a sign indicating a one-way street unless it is conventionally agreed to be so. The same applies to all the preceding examples. Representations, in everyday life, seem to be things intended to stand for other things. Whatever "standing for" amounts to is not important. What is important is that intentions enter in the nature of everyday representations but not in the nature of phenomenal experiences. It is in any case definitely not a tenet of the representationalist view that experiences are intended to stand for other things, on any reasonable interpretation of "stand for".

One might think that we could successfully pin down the relevant kind by combining examples of representations and canonical attitudes. But what is picked out by "the kind common to everyday representations and canonical attitudes" is subject to the same indeterminacy as "what is common to all canonical attitudes". As before, we can take either the most specific, the most natural, or the least specific kind. The most specific plausibly has such properties as not being phenomenal, which rules it out for the representationalist's purposes. The least specific is Entity or State. As before, there are also many equally interesting natural kinds exemplified by all the relevant examples—more or less every philosophical theory of representation describes one.

Let us take stock. In this section I have discussed two traditional ways of explicating the notion of intentionality or representation which is central to representationalism as typically formulated: through the metaphor of aboutness, and through examples. Explanations of these kinds might be sufficient for some purposes. They might be sufficient to convey a rough and ready sense of what various phenomena in the neighborhood of intentionality have in common. But our goal is to flesh out a representational theory of consciousness which can provide a framework in which

to formulate clear and precise hypotheses about consciousness. Traditional explications of the notion of intentionality either a) leave it too vague or too obscure for this purpose or b) yield an understanding of representationalism which makes it implausible.

1.2 Beyond aboutness and ostension

The preceding points are largely unoriginal. I would say that most theorists are aware at least in the back of their minds of the considerable vagueness of "aboutness", "representation", "intentionality" and derivative terminology (e.g. "content") as typically introduced. A number of representationalists attempt to explain their views in more precise terms. For example, Byrne (2001) appears to take it as axiomatic that the content of an experience is "the way the world perceptually seems to [its] subject". Others, such as Chalmers (2004), Siegel (2005b; 2006; 2007), and Siewert (1998), assimilate the contents of experiences to accuracy, veridicality, or satisfaction conditions.[6] Of course, there are also theorists who understand the representationalist view in terms of information-theoretic notions of representation, e.g., Dretske (1995), Lycan (1996a), and Tye (1995; 2000).

At this stage, it might seem that there is no adequately clear common element that unifies all representational theories, and that we would do best to settle for one of the preceding approaches to representationalism instead of trying to unify them. I will later argue that there is a viable theory which unifies these approaches, but for now I want to do something else: I want to briefly indicate some difficulties with these approaches and a few others. This will help motivate the radical departure from prior formulations of representationalism I am about to propose. Another reason why I want to discuss these difficulties is that similar issues have repeatedly been presented as problems for representationalism generally; I want to highlight

[6]Siewert does not identify contents with accuracy conditions but holds that having accuracy conditions is sufficient to have content.

the fact that they arise only on particular interpretations of the view.

Four relatively clear ways of cashing out the idea that experiences "have content" have gained currency in the literature:

1. Experiences carry information about the world.

   e.g. Dretske (1995), Lycan (1996a), Tye (1995; 2000)

2. Experiences have accuracy, veridicality, or success conditions.

   e.g. Chalmers (2004), Siegel (2005b; 2006; 2007), Siewert (1998)

3. Experiences make things seem, look, or appear a certain way.

   e.g. Byrne (2001), Byrne & Hilbert (2003), McGinn (1989;1988)

4. Experiences are acquisitions of potential beliefs or invitations to belief.

   e.g. Armstrong (1968), Pitcher (1971), Jackson (2004)

There is another, less common approach adopted by Crane (2003) and Pautz (2007; 2009). Crane and Pautz characterize the intentional contents of experiences in part in terms of the latter's relational structure.[7] This is the approach I favor, and I will return to it later. Aside from it, all extant accounts of the intentionality of experience which go beyond the unanalyzed notion of intentionality or representation fall under the heading of one of the preceding views.[8] I now want to briefly outline the main problems each of these views faces.

Representationalism says (among other things) that experience is necessarily representational. If we understand representation in information-theoretic terms, we obtain a view which entails that experiences necessarily have informational properties. But this consequence appears dubious. One way to bring out its dubiousness

[7] Crane is less explicit than Pautz about this, but this seems to me to be the gist of his solution to the puzzle we are facing.

[8] I don't count topic-neutral analyses of the kind suggested by Smart (1959) and Lycan (1987) as providing accounts of the intentionality of experience. On topic-neutral accounts, an experience of red is (roughly) a state of the kind one is normally in when one perceives something red. The problem with analyses of this kind is that there are very many states one is normally in when one perceives something red or any other kind of thing.

is with a by now familiar combination of the brain-in-a-vat and Swampman scenarios. Imagine an alien brain which pops into existence in empty space somewhere, has some color experiences, then disappears. The brain in question has no lineage (hence no biological function), no sensory equipment of any kind, and is not at all similar to a human brain. By construction, this brain's experiences would not have informational properties of any kind that has been invoked in a theory of consciousness. If such a brain is possible (and it really seems to be), phenomenal states are not informational states of any relevant kind. More problems for this type of representationalism are raised by Block (1990; 1998; 2003), Ellis (2007), Gray (2003), Kirk (1996), Macpherson (1999; 2003; 2005; 2006), Nickel (2006), Pautz (2006), and Wager (1999).

One problem which has not been discussed much but seems critical to me is that no information-theoretic account of intentionality seems to achieve the sharpness required to plausibly characterize the nature of experience. When consciousness is assimilated to an informational process, one gets the impression that it is being described with the precision of a blueprint for a computer chip. Not so. The notion of information which is at play in such theories as Dretske's and Tye's is not one of the engineering notions which go by this label but something far more nebulous.

According to Dretske, a state represents a given state of affairs just in case it has the function of indicating this state of affairs. In his 1995 and 1988 books, Dretske refers back to his earlier work for a precise definition of indication. But the notion has been shifting through his earlier work. In his 1981 book, he gives a very demanding, highly precise definition of information according to which a state carries the information that P just in case the probability of P given that the state obtains is 1. If indication is understood this way, plausibly no brain state has the function of indicating an external state of affairs. The problem is that a state can only acquire its function by performing it, and plausibly no brain state has ever stood in such a perfect correlation to external states of affairs. Problems

of this type led Dretske to weaken his account of the kind of indication relevant to representation (see Godfrey-Smith 1992). In his 1986 article, he begins using the terms "indication" and "natural sign" rather than "information" and suggests a less demanding account of the notion: he describes natural signs as "more or less reliable indicators" (p. 18). Indication no longer seems to require a perfect correlation. But what strength must the correlation have, then? Dretske never explains. Should we say that an experience represents an F just in case it has the function of correlating with Fs with probability 0.8654?

There is also the issue that Dretske does not tell us precisely how states acquire natural functions. On most accounts, having a natural function is a matter of having been useful in a certain way during evolution. We can try to give a precise characterization of how useful a state must have been to acquire a function, or we can leave the notion vague. We again seem to face a choice between vagueness and implausible precision.

Parallel remarks apply to Tye's account of content in terms of covariation in "optimal conditions". The problem with both Dretske's and Tye's suggestions is a dilemma: if they are made more precise, they will likely seem utterly arbitrary; if they are left vague, they cannot be literally, entirely correct and complete, because consciousness is not vague.

I now turn to the second approach. Let us use the term "success condition" in such a way that success conditions include all potentially relevant kinds of condition: accuracy conditions, veridicality conditions, satisfaction conditions, and what we would normally call "success conditions", whatever these are.

The claim that experiences essentially have success conditions has been criticized from many directions. Opponents include Brewer (2006; 2007; 2008), Travis (2004), Smith (2008), Biggs (2009), and Pautz (2009). Smith bases his case on blurry experiences. Brewer claims that the Müller-Lyer illusion shows the possibility of experience*s* without accuracy conditions because there is no way the world

could be that would make such an experience accurate. Biggs offers the fictitious example of streams of randomly occurring experiences (visual experiences similar to white noise on TV), which, he argues, would have no accuracy or veridicality conditions. Travis gives an argument which turns on the idea that the accuracy or veridicality conditions of an experience, if any, should be reflected in how things look to one when undergoing it. Pautz distinguishes two accuracy conceptions of phenomenal content: one grounded in examples, and one grounded in an explicit definition. He then argues that both conceptions of phenomenal content fail to capture the intended notion of content because they trivialize central questions about phenomenal content.

My view is similar to Pautz' in outline. The main problem with success conditions takes the form of a dilemma. Either we understand the relevant conditions on the model of familiar entities which have them, or we do not. In the former case, experiences don't seem to have success conditions. In the latter case, it is most unclear what success conditions might be. I will briefly discuss each branch of this dilemma.

We can ask first if experiences have belief-like veridicality or accuracy conditions, or desire-like satisfaction conditions, that is, if they have the potential for being veridical or accurate like beliefs, or satisfied like desires. The main challenge for this proposal is to account for idle episodes of sensory imagination, including idle episodes of visual and aural imagery (as when one talks to oneself). These kinds of imagery involve experiences which must be accounted for by representationalism, because representationalism is a general theory of consciousness which must apply to all experiences. To illustrate the problem, take the jingle that is currently stuck in my head (the "Zoom Zoom Zoom" jingle of Mazda's television ads). I am not, to use Siewert's (1998) phraseology, assessable for accuracy in virtue of experiencing the jingle. My experience does not seem to aim at truth in the same way as beliefs. It also does not seem to aim at satisfaction in the same way as desires:

I would not be satisfied if the jingle started playing. My jingle experience does not even seem to have the direction of fit of belief or desire. Of course, the idea that experiences are belief-like or desire-like is vague and one could insist that there is a similarity. But anything is similar to anything. Whatever jingle experiences have in common with beliefs and desires is not sufficiently salient that we can leave the matter at this. Saying that jingle experiences have success conditions like beliefs and desires does not unambiguously pinpoint an interesting feature of experiences.[9]

An alternative is to understand the relevant success conditions in terms of biological function. But we have already found a major difficulty with the claim that experiences essentially have biological functions, so we can set this possibility aside (this is also not how proponents of the success condition view see it).

The only other way of explicating the relevant success conditions is by comparison with things lay people would sometimes describe as "successful", "accurate" or "satisfactory". Examples include job applications, military interventions, gun shots, portraits, measurements, theories, descriptions, chemical analyses, directions, projects, etc. All of these have success conditions, whether accuracy conditions or success conditions of some other kind. However, they all have their success conditions in virtue of certain people having certain intentions or expectations. Experiences (jingle experiences, at least) are not like that, so we cannot explain the success conditions of experiences by comparison with these mundane success conditions.

We have now reached the second horn of the dilemma. If experiences have success conditions, it is not in a sense which can be immediately grasped through

[9]David Chalmers (private correspondence) suggests that one could give different accounts of the contents of sensory imagination and perceptual experiences. More specifically, he suggests that part of the difference between sensory imagination and perceptual experience could be that only the latter's contents are proposition-like. The contents of sensory imagination could be akin to the contents of predicates, for example.

This proposal is compatible with sensory imagination not having success conditions, but it does not present a solution to our problem: our problem is to explain what it is for an experience to have content in such a way that representationalism turns out true. If an experience's having content is its having success conditions and experiences of sensory imagination do not have such conditions, then not all experiences have content and representationalism is false.

examples of other things that have success conditions: it is not in the everyday sense, in the biological sense, or in the sense in which philosophers say that beliefs and desires have success conditions. In what sense? What is the relevant technical meaning given to "success condition"? Some (e.g. Chalmers 2006) give examples of successful and unsuccessful experiences to help fix ideas, but such examples leave the relevant concept highly indeterminate (Pautz 2009: 488-9).

As far as I can tell, proponents of the success condition view have not attempted to clarify the relevant notion of a success condition otherwise than by pointing to examples of the kinds we have already considered or by assimilating success conditions to contents of some other type under consideration. Note in particular that Siegel's (2005a) suggestion that the relevant success conditions correspond to what is "conveyed" to the subject of an experience seems to equate the success condition view with the view that the content of an experience is what it inclines or invites one to believe (which is discussed below).

One is tempted to say that typical jingle experiences have accuracy conditions (but maybe not success conditions in the other senses considered so far) because they can be *compared* or *matched* with the world in some way. But anything can be matched with anything. I challenge proponents of such attenuated accuracy conditions to give an account of them which does not confer accuracy conditions to the rug in my office. Like experiences, my rug can be compared or matched with the world. There is a function F which accepts rugs as input and outputs propositions. Its output for my rug is the proposition that a certain jingle is playing. I can compare the F-value of my rug with the way the world is: its F-value can be accurate or inaccurate in the sense that it can be a proposition that is true or false. We are inclined to compare experiences with the way the world is with respect to what they are experiences "of". Rugs can be compared with the way the world is with respect to their F-values. What is the relevant difference?

It is tempting to say that comparing what an experience is "of" with the way

the world is is less arbitrary than comparing a rug's F-value with the way the world is. I can think of only one way of cashing out this idea. The idea must be that the relevant difference between my rug and my jingle experience is that it is an essential part of the latter that it is "of" the tune in question, while it is not an essential part of my rug that its F-value is this tune. But this only moves the bump in the rug.[10] For consider the state S of owning a rug whose F-value is the proposition that the Mazda jingle is being aired. S essentially involves standing in a certain relation to the same proposition or state of affairs my experience is "of". Why doesn't my rug have the same accuracy conditions as my experience? At this stage I don't know how a proponent of accuracy conditions might further develop the proposal. One wants to say that S lacks the right kind of aboutness or directedness, but that is the very notion we are trying to clarify.

Let us now consider Armstrong's, Pitcher's and Jackson's suggestions to the effect that the contents of experiences should be understood in terms of relations to beliefs. Armstrong and Pitcher suggest that an experience has a given content P if it is an acquisition of a potential belief that P. Jackson proposes that experiences are invitations to believe propositions. The same example which makes trouble for accuracy conditions makes trouble for this general approach: the jingle in my head does not seem to even begin to induce a belief in me. I do not feel inclined or invited to believe that the tune is playing.

The problem is easier to see if we keep in mind that believing or desiring requires robust cognitive dispositions which reflect what one believes or desires (as argued above). It does not seem that experiences always come with sufficiently fine-grained dispositions. Take for example the total, unified visual experience I have of my office at this time (call it o). Millions of different colors and even more small surfaces figure in the content of o. Now take an experience o' which is just like

[10]Besides, it is an essential part of my rug that it has the proposition in question as F-value if all necessary properties are essential properties. A representationalist ought to be uncomfortable hanging his or her theory on the metaphysical thesis that some necessary properties are not essential properties.

o except that it presents the top-left most perceptible point on the copy of Hume's *Treatise* two meters away from me as being of a slightly different shade of purple. I doubt that one's cognitive dispositions with regard to experiences such as o and o' are always different. Yet they would have to be for a belief-grounded representational theory to account for all differences in phenomenal character between them. In the same vein, cases of change blindness with respect to striking features of objects of attention seem to show that some differences in phenomenology (even quite stark ones) sometimes fail to be reflected in the connections phenomenal states stand in with cognitive states (see Levin et al 2002 for some relevant cases).

The Mazda jingle and the coarse grain of inferential links between experiences and beliefs are also a problem for the view that experiences make the world seem, appear, or look a certain way to one. The world's seeming, appearing, or looking a certain way can be understood phenomenally, comparatively, or epistemically. Understood either phenomenally or comparatively, to say that the world seems, appears, or looks a certain way is just to say that one is having experiences of a certain kind. In the comparative case, it is to say something roughly along the lines of "this causes in me experiences similar to those an X would (probably) cause". Understood this way or phenomenally, the claim that experiences make the world seem, appear, or look a certain way has no interesting content. Understood epistemically, the world's seeming, appearing, or looking a certain way is a matter of one's having visual evidence for a claim, or perhaps an experience which inclines to believe a certain claim. We saw already that not all experiences bear such epistemic connections to beliefs: my jingle experience does not. The epistemic connections between experience and belief are in any case too coarse-grained to reflect all possible variations in phenomenology. So the phenomenal character of experience cannot be fully explained solely by reference to how the world seems, looks, or appears to one in the epistemic sense. Travis (2004) and Pautz (2009) offer more extensive cases against representationalism understood in terms of seemings, appearances, or

looks.

I hope that the brief remarks in this section will have sufficed to convey the difficulties I want to avoid by setting aside the preceding approaches. It is common to conclude that representationalism is false on the ground that experiences do not satisfy any or some of these accounts, but this line of reasoning ignores the fact that each of these proposals is an attempt at precisifying the vague idea that experience is a species of intentionality or representation. If this idea could be precisified differently, representationalism might be saved. I will now describe the approach I favor.

1.3 Virtualism

To isolate the core representationalist theory, and so to find out how a representationalist should explicate the notion of content, we need only ask ourselves in what respects phenomenal states have to be like canonical intentional states or representations for the main consequences of the theory to follow. In introduction I said that the two main applications of the theory had to do with the regimentation problem and the problem of perception. It has other applications, which will be discussed in the next chapter and throughout this thesis, but we can already begin to see that the theory's applications do not seem to hinge on considerations about representation or intentionality *per se*. The two applications discussed so far depend on two points: a) that phenomenal states have a certain relational or quasi-relational structure involving on one end proposition-like entities (this is what helps with the regimentation problem); b) that phenomenal states have a certain fallibility or independence from their objects (this is what helps with the problem of perception). In the next chapter, I will argue that all the main applications of the theory turn on these two points and nothing else. Assuming for the moment that this is correct, the following statement would seem to capture the essence of representationalism:

Definition 1 There is a nonfactive relation R such that, for every phenomenal state *s*, there is some proposition P such that *s* = standing in R to P.

A relation R is factive just in case something's standing in R to some proposition P entails that P is true. So a nonfactive relation is a relation that something can stand in to a false proposition if something can stand in it to any proposition at all.

Definition 1 makes reference to propositions, but this should not be read too narrowly. Consistently with the requirements on representationalism uncovered so far, I mean to count as propositions all proposition-like things. I write "propositions" because one tires quickly of "proposition-like things". I also happen to think that the term "proposition" is used sufficiently loosely in the literature that it already can encompass almost all proposition-like things.

What is a proposition-like thing? Perhaps the most familiar examples of propositions are sets of possible worlds. In addition to the various set-theoretic constructs which have been labeled "propositions", I want to include in the category of proposition-like things world-properties (properties of whole possible worlds), possible situations (in the sense of Barwise and Perry (1983)), Kim-style events (instantiations of properties), and states of affairs. I count as propositions all entities which can easily and systematically be described using statements in higher-order predicate logic, because what matters for the applications of representationalism is that the contents of experiences be readily and systematically describable using such a formalism.

I am inclined to think that the best candidate relata of experiences are states of affairs of some kind or other, but I will for the most part remain neutral on this here. I will also remain largely neutral on what a state of affairs is. However, I take it as axiomatic that there is a difference between facts and states of affairs: states of affairs are not facts but things of the kind which are facts when they obtain. One might want to assimilate states of affairs thus understood to ways the world could be, states of the world (properties of the world), arrangements of objects and properties,

abstract structures composed of universals, or some other kind of entity—I leave all these options open.

Definition 1 is a good start, but it has two significant flaws. First, it does not guarantee that phenomenal states are fallible. The reason is that relation R could have instances which are not constitutive of phenomenal states. Since nonfactivity at most requires that one can stand in R to some false proposition, R could be fallible with respect to its non-phenomenal instances while being infallible with respect to its phenomenal instances. But we want to capture the fallibility of relation R with respect to phenomenal states in order to block the argument from illusion for sense data (more on this topic in section 2.1).

Another problem is that definition 1 does not guarantee that some phenomenal states at least can be both veridical and non-veridical (in different circumstances). In other words, it is consistent with definition 1 that every phenomenal state is either necessarily veridical or necessarily nonveridical. This means that the view captured by definition 1 is not strictly speaking incompatible with representationalism's main competitor (disjunctivism; see below for some possible definitions of disjunctivism).

We can improve on definition 1 by invoking the tailor-made notion of *virtuality*.

Virtuality A relation R is virtual with respect to a set S of its state instances iff there is a state *s* in S which is a state of standing in R to some *x* and which is such that it is metaphysically possible for *s* to obtain whether *x* obtains or not.

By "state instances" I mean states (or properties) which consist in standing in the relation to a given entity. For example, *standing in the next-to relation to my desk* is a state instance of the next-to relation. A state instance of a relation R is a monadic property formed by assigning a value to all but one of the arguments of R. I use the term "obtain" to remain neutral regarding the precise nature of the relata of the relation. We can say that propositions obtain when they are true, that properties

obtain when they are instantiated, and that individuals obtain when they exist.

Given this notion of virtuality, we can restate the core representationalist thesis as follows:

Definition 2 There is a relation R such that 1) for every phenomenal state *s*, there is some proposition P such that *s* = standing in R to P; 2) R is virtual with respect to phenomenal states.

Definition 2 solves the preceding problems, but it is too strong in one respect. Consider these three phenomenal states:

(1.5) Experiencing a red square floating above the ground

(1.6) Experiencing some color or other

(1.7) Having an experience

In the first case, there is an obvious candidate proposition for being the content of the experience: that there is a red square floating above the ground. But what propositions could satisfy definition 2 in the two other cases? These don't seem to be states in which one experiences a complete state of affairs, yet we should arguably count them as phenomenal states.

What distinguishes the second and third states from the first is that they are *derivative* phenomenal states, while the first is a non-derivative or *basic* phenomenal state.

Derivative phenomenal state A state which consists in being in one of a given set of phenomenal states distinct from itself.

Basic phenomenal state A phenomenal state which is not derivative.

For example, state 1.6 consists in having one of the possible color experiences, where color experiences are experiences of specific colors. It is not itself a color experience. It is therefore derivative. It is arguable that most of the phenomenal

states we discuss in academic contexts are derivative. Take the state of experiencing red. When we speak of the experience of red, we generally mean to refer to a kind of experience, namely, experiences which are of red things among other things. We count experiencing red as a phenomenal state, but the only sense in which there is "something it's like" to experience red is that when one experiences red, one experiences certain things having certain properties, one of which at least is the property of being red. I will return to the question of which states are basic and which are derivative in chapter 3. For now, we need only keep in mind that there is a distinction, and that derivative phenomenal states are mere states of having phenomenal states of a certain kind.

A theory of consciousness which explains basic phenomenal states would automatically explain derivative phenomenal states, so we only need a theory of basic phenomenal states. Although this is seldom made explicit, most theories of consciousness are aimed primarily at explaining basic phenomenal states. Representationalism makes no exception.

We can now state the core representationalist claim, which I am going to refer to as *virtualism*.

Virtualism There is a relation R such that: 1) For any basic phenomenal state *s*, there is a proposition P such that *s* = standing in R to P; 2) R is virtual with respect to basic phenomenal states.

The idea remains roughly that phenomenal states are like propositional attitudes in that they are nonfactive relations to propositions.

Note that virtualism makes no reference to intentional or representational contents. On this view, the content of an experience is simply the proposition one is related to in it.

Earlier I said that the approach I wanted to suggest is similar to those adopted by Crane (2003) and Pautz (2007;2009). It is similar to Crane's and Pautz' approaches in that the latter also explain the notion of phenomenal content in part by

reference to the relational structure of experience. But there are significant differences. Crane describes the relations which are constitutive of intentional states as "intentional modes", and he explicates the notion of an intentional mode primarily through the example of propositional attitudes. This is where the approach I am proposing differ from Crane's. As I tried to show earlier, we cannot make a suitable notion of intentionality sufficiently precise by pointing to propositional attitudes as examples. The problem is that it is unclear how to generalize from the canonical propositional attitudes, while we don't want to say that the only intentional modes are those involved in canonical propositional attitudes.

Pautz completely reverses the usual order of explanation between intentionality and representationalism. Rather than explicating representationalism in terms of some independently given notion of intentionality or intentional content, he explicates the notion of an experience having content in terms of representationalism. Virtualism is inspired from Pautz' approach in this regard. Simplified a little, Pautz' intentionalism says that there is a relation R such that phenomenal states are identical to states of standing in R to proposition-like entities.[11] Virtualism differs from this view in asserting that relation R is a virtual relation and not any old relation. My main quibble with Pautz' intentionalism is that it does not guarantee the fallibility of experience.[12]

A noticeable shortcoming of virtualism and Pautz' intentionalism, compared to Crane's view at least, is that the former give us no clue as to what the contents of

[11] Pautz' statement of his intentionalist theory:

> There is a distinctive relation R such that for every experiential property of the form having an experience with minimal phenomenal character K, there is some intentional content c, such that the property of the form having an experience with minimal phenomenal character K is identical with the property bearing R to c (or else there is some type of intentional content T such that the experiential property is identical with standing in R to some singular content or other of type T). (2007: 497)

[12] Pautz' qualification that the relevant contents are "intentional" (see previous footnote) might seem to play the same role as my qualification that the relevant relation is virtual. However, the point of Pautz' approach, as I understand it, is to avoid having to define intentionality, contents and cognate notions independently of representationalism. It is only without this qualification that Pautz' definition solves our definition problem.

particular experiences are or what relation R is. Virtualism is incomplete in this respect, and its applications slightly restricted because of this. I will state a more robust theory (virtualism+; p. 141) when the need arises. This theory will fall out of my case for the basic virtualist theory stated above.

Earlier I said that I am inclined to think of the propositions (i.e. proposition-like things) we are related to in experience as states of affairs. On this understanding of virtualism, it arguably commits one to there being non-obtaining states of affairs: if one is related to a state of affairs which does not obtain, it follows by existential generalization that there is a state of affairs which does not obtain. One might think that this is a heavy ontological burden for a mere theory of consciousness.

This objection raises a host of issues in ontology which I cannot discuss in detail here—there are considerations against virtualism which are more specific to it I want to concentrate on. However, two points should help alleviate this concern.

Firstly, ontology is hostage to its applications. There is no prospect of directly determining whether any given abstract object exists or not. The only sensible way of going about determining whether any given abstract object exists is to ask whether we need to posit its existence to explain other things. So we should accept the existence of non-obtaining states of affairs if we find ourselves having to talk about them in our best theory of the world. Since a theory of consciousness would be an application of major importance for non-obtaining states of affairs, it would be backward to reject virtualism on the ground that there are no non-obtaining states of affairs. Rather, we should first ask whether virtualism and other relevant theories outside of ontology are plausible independently of such considerations, then draw ontological conclusions from these theories.

Secondly (and relatedly), it does not seem implausible that virtualism shares its ontological commitments with the theory of propositional attitudes. The reason I, at least, am inclined to say that experiences are relations to states of affairs rather than propositions of a more abstract kind is that this seems to me to be the most

phenomenologically accurate view. A parallel observation applies to propositional attitudes. It seems to fly in the face of common sense to say that we believe and desire set theoretic constructs such as functions from worlds to truth values or sets of possible worlds. One cannot intelligibly say "I believe the set of possible worlds S" or "I believe function F".[13] Propositions of the state of affairs kind seem to be the most likely candidates for being the objects of beliefs and desires. If that is correct, then there is good reason independently of virtualism to think that there are non-obtaining states of affairs.

1.4 Other representationalisms

There is one respect in which virtualism is likely too committal to reflect the view nearly all representationalists endorse: it implies that no two distinct phenomenal states can have the same content, while several formulations of representationalism leave this open. For example, Chalmers (2004), Crane (2003), John (2005), and Lycan (1987; 1996a) appear willing to allow that phenomenally distinct experiences in different sensory modalities have the same content. Virtualism does not allow this, but it could be weakened for this purpose.

Sensory virtualism (SV) For any sensory modality M, there is a relation R such that: 1) For any basic phenomenal state *s* in M, there is some proposition P such that *s* = standing in R to P; 2) R is virtual with respect to basic phenomenal states.

SV makes room for the kinds of possibility that concern the aforementioned authors, but it has the major shortcoming of covering only phenomenal states which can be associated with sensory modalities. It is unclear how to generalize it to all phenomenal states. I will return to this matter in in chapter 5.

There is also another view it is useful to distinguish:

[13] See the substitution problem discussed in McGrath (2008)

Weak virtualism (WV) For every basic phenomenal state *s*, there is a virtual relation R and a proposition P such that: 1) *s* = standing in R to P; 2) R is virtual with respect to basic phenomenal states.

Weak virtualism allows any number of additional ingredients (including intrinsic qualia) to play any role whatsoever in the determination of an experience's phenomenal character beyond its content—whatever they are, they can be built into R. As such, it is much weaker than SV and virtualism (though stronger than SV in that it covers all phenomenal states).

It may be argued that the sensory view, or even the weak view, is the one which ought to be described as the core representationalist theory. They can handle a number of the same explanatory tasks, including the problem of perception. However, they seem much less satisfactory than virtualism as far as the regimentation problem goes. We will see in chapter 5 how much worse than the pure virtualist theory the sensory and weak views are at dealing with the regimentation problem. I describe virtualism as the core representationalist view because I believe that the solution to the regimentation problem it enables is a pivotal motivation for representationalism generally.

Virtualism is compatible with nearly all representational theories I am aware of, including the sensory and weak views just introduced (these theories do not exclude the possibility that the relations which satisfy them are all identical to the relation posited by virtualism). Both reductive and ontologically neutral theories are compatible with virtualism: Armstrong (1968), Pitcher (1971), Dretske (1995), Lycan (1996), Rey (1998), Harman (1990), Tye (1995), Anscombe (1965), Chalmers (2004), Byrne (2001), Crane (2003), and Jackson (2004) could all happily endorse virtualism. Reductive theories can be regarded as providing reductive accounts of the relation R posited by virtualism. For example, Tye's theory that experiences are nonconceptual, poised states which covary with their objects in optimal conditions can be regarded as virtualism combined with a reductive account of relation R. On

his view, R would be *being in a state which covaries with* x *in optimal conditions while being nonconceptual and poised*. Virtualism can also accommodate some higher-order and self-representational theories, e.g. those advocated by Rosenthal (2002) and Kriegel (2003). For instance, one can take relation R to be the relation of *representing* x *while having a higher-order thought about the fact that one is representing x*. R could not be *representing that one instantiates a quale x while one is instantiating x* (this relation is not virtual, and its relata are not propositions), so some higher-order theories are ruled out. Still, virtualism and its weaker variants seem to capture the essence of centrally representational theories of consciousness.

1.5 Opposing views

Proponents of adverbialism and the intrinsic qualia theory (e.g., Chisholm 1948, Ducasse 1942, Block 2003) oppose the idea that experience has an act-object structure—that it can be decomposed into a relation common to all experiences and objects which vary from one experience type to another. This idea is explicitly captured by virtualism, though virtualism posits that the objects of experiences are propositions, not the phenomenal individuals posited by sense-datum theorists. Virtualism also seems to capture what sense-datum theorists (e.g., Jackson 1977, O'Shaughnessy 1980, Robinson 1994) oppose. As we will see in the next chapter, the starting point for this theory is precisely that experiencing an F is not a virtual relation. Sense-datum theorists often put the point by saying that the relevant relation is a kind of "acquaintance".

The relation between disjunctivism and virtualism is less straightforward.

Two characteristics are shared by all positions which fall under the heading of disjunctivism. The first is a claim (or cluster of claims) disjunctivists label "naive realism". To a first approximation, naive realism asserts that there is a strong connection between veridical experiences and external entities. This connection has

been conveyed using such expressions as "direct contact", "immediate presence", and "openness to the world". The connection is taken by disjunctivists to imply that veridical and non-veridical experiences are importantly different—that they have no "highest common factor". This last claim is the second tenet of disjunctivism.

On the face of it, virtualism and disjunctivism seem incompatible, because virtualism posits a common factor between veridical and non-veridical experiences. William Fish's (2009) characterization of the naive realist view which is central to disjunctivism appears to make the latter incompatible with virtualism:

> The distinctive feature of naïve realism lies in the claim that, when we see the world [i.e. in veridical perception], the subject is acquainted with the elements of the presentational character—the mind-independent objects and their features—where 'acquaintance' names an irreducible mental relation that the subject can only stand in to objects that exist and features that are instantiated in the part of the environment at which the subject is looking. [...] Why is it like that to have that experience? Because in having the experience, the subject is acquainted with thus-and-such objects and their properties. This acquaintance property can therefore be identified with the experience's phenomenal character. (Fish 2009: 14-5; comment in square bracket is mine)

Going by Fish's account, naive realism appears to imply that the phenomenal states which are tokened in veridical experience are states of standing in a factive relation to external states of affairs. It is not clear to me that naive realism as traditionally understood has this implication, so I will refer to this claim as *factualism* to avoid confusion.

Factualism The phenomenal states instantiated in veridical perceptual experience are states of standing in a factive relation to states of affairs.[14]

[14] A factualist would probably want to restrict this claim to basic phenomenal states as I have done above for virtualism.

A factualist like Fish must deny that the same phenomenal states can be tokened in veridical and non-veridical experience—to deny a common phenomenal factor to veridical and non-veridical experience. This is one natural understanding of the disjunctivist position. This understanding of the position makes it incompatible with virtualism, because virtualism entails that some phenomenal states can be instantiated either in veridical or non-veridical conditions.

There are at least two alternative understandings of naive realism which do not carry any commitments to factualism:

Objectivism The phenomenal states instantiated in non-hallucinatory perceptual experience are states of standing in an acquaintance relation to ordinary objects.

Particularism The phenomenal states instantiated in non-hallucinatory perceptual experience are states of standing in a relation to states of affairs involving ordinary objects.

The main difference between factualism and these two views is that factualism implies that illusory experiences instantiate different phenomenal states than veridical experiences. On the other hand, all three views imply that veridical and hallucinatory experiences instantiate different phenomenal states (assuming there are no non-existent ordinary objects).

Objectivism is no more compatible with virtualism than factualism is, but for different reasons. The incompatibility comes from the fact that virtualism implies that the phenomenal states instantiated in veridical perception are relations to propositions and not particulars. However, it is worth noting that virtualism is compatible with a number of claims which are commonly taken to support objectivism. In particular, it is compatible with the claim that we experience ordinary objects, because we can do this as part of experiencing states of affairs. Virtualism is compatible with the claim that we experience ordinary objects even on the assumption that one

can experience an ordinary object only if the object exists and one is suitably related to it. I will discuss these matters at some length in section 7.2.

In contrast with objectivism, particularism is entirely compatible with virtualism, because the latter at most requires that illusory and veridical experiences share phenomenal character, while particularism at most implies that hallucinatory and veridical or illusory experiences never instantiate the same phenomenal states. Virtualism might seem a little too weak because of this, but building the negation of particularism into the theory would not add to its explanatory power as far as its core applications are concerned. One could strenghen virtualism to say that the propositions one is related to in experience never involve particulars, but I prefer not to do this because this would raise subtle phenomenological issues which are orthogonal to my primary aims. I will nevertheless return to the question of object-involving contents in section 7.2.

Factualism seems to me to be the most common understanding of the core tenet of naive realism. For example, Campbell (2002: 116) appears to be a factualist, and Hellie's (2007; 2006) exposition of naive realism is along the same lines as Fish's. The picture painted in broad strokes by McDowell (1982; 1986; 1994) also seems to me to have something like Fish's factualism as a central component, though there are dissenting interpretations.[15]

Martin (2004) can easily seem to endorse factualism, but there is considerable room for alternate interpretations of his position. Take for example this passage:

> The Naïve Realist, however, claims that our sense experience of the world is, at least in part, non-representational. Some of the objects of perception—the concrete individuals, their properties, the events these partake in—are constituents of the experience. No experience like this,

[15]Byrne & Logue (2008) suggest that McDowell is only a disjunctivist about the epistemic status of experiences. There is certainly room for such an interpretation. However, we will see in chapter 8 that McDowell's arguments lend themselves easily to readings on which they support a Fish-type naive realism. It seems to me that this is how they are normally understood. Since this is the only interpretation of McDowell's work which makes contact with virtualism and it seems to be the most common interpretation, I will adopt it for present purposes.

> no experience of fundamentally the same kind, could have occurred had no appropriate candidate for awareness existed. In this, sense perception contrasts with imagining and thought. For one can certainly imagine objects in their absence, so the mind's direction on an object does not require that it actually exist when one imagines. The same is true, arguably, of thought—we think of objects which in fact do not exist as well as thinking of the existent. The Naïve Realist insists that sensing is not like this, and in that respect the Naïve Realist finds common ground with the Sense-Datum tradition, or what more broadly I will label Subjectivism. For Subjectivists have long insisted that what is distinctive of sensing as opposed to thinking is that one really cannot sense in the absence of an object of sensing. (2004: 39)

The second sentence of this passage suggests that Martin is using the term "object" to refer to what an experience is directed at (as opposed to the material object one perceives). He also seems to take the primary objects of experiences to be complex events or states of affairs, not ordinary objects. Since he claims that the directedness of experiences toward their objects is not a kind of representation, he seems to endorse factualism. After all, the only thing that is clear about representation is that it is not a factive relation; if Martin denies that the relation constitutive of experience is a kind of representation, he must think that it is factive.

Having said this, there is room for alternate readings of Martin's position. In particular, the part of the above passage which is most characteristic of what Martin says elsewhere is the claim that "no experience like this, no experience of fundamentally the same kind, could have occurred had no appropriate candidate for awareness existed". This claim does not clearly imply factualism (or particularism or objectivism). The key question is what "fundamentally the same kind" means. All Martin says in guise of clarification is that he does not mean "fundamental" in the sense in which particle physics is concerned with fundamental particles, and

that we should not try to determine whether two things are fundamentally of the same kind "merely through appeal to an inclination on our part to describe some things as similar and others as different" (Martin 2006, footnote 9). In other words, we should allow that things are not fundamentally of the same kind even when they strike us as similar. But what is it that we are allowing? What is it for two things to be fundamentally of the same kind? This question is crucial to the relation between Martin's disjunctivism and virtualism: his position, if defined by the preceding statement, would appear to be incompatible with virtualism just in case sameness of phenomenal character implies sameness of fundamental kind in his sense. If phenomenal sameness implies fundamental sameness, virtualism is incompatible with Martin's position, because virtualism implies that the same phenomenal states can occur independently of the states of affairs they relate us to. If the phenomenal state instantiated in an experience is its phenomenal kind, then virtualism implies that veridical and non-veridical experiences of the same fundamental kind are possible. However, if phenomenal sameness does not imply fundamental sameness, then virtualism seems to be compatible with Martin's position.

Does phenomenal sameness imply fundamental sameness in Martin's sense? One cannot say for sure, because Martin does not sufficiently explain what he means by "fundamental". However, it does not seem improbable that the phenomenal state instantiated in an experience should count as its fundamental kind in Martin's sense. The reason is that the phenomenal state it instantiates is by far the most salient and interesting property which enters into its nature. But we cannot say for sure given what Martin discloses about his unconventional use of "fundamental".

In brief, Martin's statement that the directedness of experiences is not representational suggests that he endorses factualism. The "fundamental kind" formulation of his position is less clear on this matter, but it can reasonably be taken to conflict with virtualism.

I should stress that it is not only the virtualist approach to representationalism

whose relation to Martin's position is unclear. Even Dretske's and Tye's views could conceivably be compatible with Martin's claim that veridical and non-veridical experiences lack a common fundamental kind: it could be that the informational-functional properties which are identical to phenomenal states on Dretske's and Tye's views are not fundamental properties by Martin's lights. This is in fact quite plausible given that Martin does not regard the epistemic, quasi-functional properties which veridical and non-veridical experiences share as sufficient for fundamental sameness (c.f. Martin 2006: 72).

Snowdon's (1980; 2005) disjunctivism is also slightly elusive. According to Snowdon (2005), "the thesis that 'disjunctivism' stands for is precisely the denial of the common visual element claim", where the common visual element claim is that perceptions and non-perceptions "are, in respect of the experiential element in them, of the same fundamental sort". Like Martin, Snowdon does not explain the notion of fundamental sameness he is appealing to. This statement of the disjunctivist position suffers from the same indeterminacy as Martin's. However, Snowdon's position in his 1980 article seems to be a disjunctivism of either the factualist or the particularist kind:

> The disjunctive picture divides what makes looks ascriptions true into two classes. In cases where there is no sighting they are made true by a state of affairs intrinsically independent of surrounding objects; but in cases of sightings the truth-conferring state of affairs involves the surrounding objects.
>
> It is this picture, rather than the claim that the actual formula given to express the disjunctive theory adequately does so, which constitutes the core-idea, on the basis of which radical alternative explanation can be given. (1980: 186)

Byrne & Logue (2008) suggest that Snowdon's position puts illusory and veridical perceptual experiences in the same category. On this reading of the preceding

passage, which Snowdon has confirmed to me is correct, his view seem to be best captured by the particularist position.

To summarize, disjunctivism combines claims of two kinds: some kind of naive realism, and a corollary claim to the effect that veridical experiences lack a certain type of commonality with certain non-veridical experiences. I identified three potential understandings of the naive realist thesis central to disjunctivism: factualism, objectivism, and particularism. The first implies that veridical and non-veridical experiences do not instantiate the same phenomenal states, while the second and third imply only that veridical and hallucinatory experiences do not instantiate the same phenomenal states. We have seen that disjunctivists generally seem committed to factualism or at least particularism. However, some theorists who explain their position in terms of "fundamental sameness" and cognate expressions might perhaps not be committed to any of the preceding positions—the textual evidence is insufficient to settle the matter.[16]

I will set aside exegetical questions regarding disjunctivism and naive realism. From now on, I will reserve the label "disjunctivism" for the most widespread view which is unquestionably in conflict with virtualism, namely, the view that the phenomenal states instantiated in veridical experience cannot be instantiated in non-veridical experience (i.e. the position motivated by factualism). There might be other kinds of disjunctivism, but they are of no concern to us here because what we want to know is whether virtualism is true or not, not whether all putative alternatives are false; we need only concern ourselves with theories that clearly conflict with virtualism.

[16]This also holds of Hinton's (1967; 1973) intricate position, which I have insufficient space to discuss here.

1.6 Summary

In the first part of this chapter, I have tried to expose a dilemma for representationalists: either define the theory using the vague terminology of "intentionality", "propositional attitudes" or "representation", or adopt one of the substitute notions of intentionality which have previously been suggested and face the difficulties these notions introduce. At times it might have seemed that I wished to reject representationalism altogether, but my ultimate aim has been constructive. I meant to make the case for a change of direction in how the representational theory of consciousness is conceived of. I also meant to show in passing that some of the problems often brought up against representationalism fail to make contact with the core of the theory (virtualism). I have concluded the chapter by exploring the main relations between virtualism and other theories. As far as I can tell, they are roughly what the relations between representationalism and other theories are generally supposed to be. In the next chapter, I will try to show that virtualism can do all the important explanatory work of which representationalism is generally supposed capable.

Chapter 2

Virtualism at work

My aim in this chapter is twofold. First, I want to address the concern that virtualism is too weak to play the important theoretical roles representationalism is generally taken to play. Representationalism has a number of implications which its proponents take to yield theoretical benefits. My primary aim in this chapter is to show that virtualism, though weaker than most other general, ecumenical representational theories, has the same relevant consequences as far as the theoretical benefits of representationalism go. My second aim is to motivate virtualism by bringing out these benefits. However, I will not argue that the theoretical motivations for representationalism discussed here really make the case for the theory. At this point, I mainly want to bring virtualism up to par with other forms of representationalism.

I will discuss four themes in connection with which the main applications of the representationalist framework arise: the problem of perception, the transparency of experience, the place of consciousness in nature, and the place of consciousness in the mind. In each case I will argue that virtualism can play the explanatory role a representational theory should play.

2.1 The problem of perception

As I noted in introduction, the central problem of perception is to provide an account of illusion and hallucination which does justice to the fact that our experience of the world is direct or immediate. One of the key motivations for representationalism is that it offers a solution to this problem.[1]

There are different ways to read the claim that experience is "immediate", and some interpretations make the claim incompatible with representationalism. Still, the most natural reading makes it compatible with representationalism. That is how I will understand it for now:

Immediacy datum One does not experience external states of affairs by being aware of other things (e.g. sense data or intrinsic qualia).

The immediacy datum captures, I think, the most common and natural reading of the claim that our experience of external states of affairs is immediate: it is immediate in the sense that experiences of external states of affairs do not consist even in part in experiencing or being otherwise aware of distinct entities which stand between us and them. Stronger readings of the immediacy point will be discussed in chapter 7.

The immediacy datum is widely held to be supported by introspection (see the transparency section below, as well as chapter 7), and denying it allegedly has dire epistemological consequences. Having said this, I am only trying to reconstruct and expose an element of the motivation for representationalism, so I will not defend the immediacy datum here.

I should also make it clear that I am not suggesting that representationalism's ability to account for illusion while preserving the immediacy datum sufficiently justifies the theory. One of its main competitor, disjunctivism, is also supposed to

[1] See Anscombe (1965), Crane (2005), Harman (1990), and Robinson's (1994) historical overview, ch 7.

achieve this. So is adverbialism, and adverbialism does not imply representationalism. If there were a compelling argument from the immediacy datum to representationalism, its key premise would have to be that no other view can satisfactorily account for the datum. I am not about to run this argument. I merely observe that a satisfactory representational theory ought to fill the theoretical role which representationalists generally take their view to play regarding the problem of perception, which is to provide an account of illusion and hallucination that is consistent with the immediacy datum.

There is a wide variety of arguments from illusion and hallucination, and other related threats to the immediacy datum (e.g. the argument from perceptual variation). Here I will use as stalking horse an argument from illusion adapted from Harman's (1990) discussion.

Argument S: the argument from illusion

1. If one perceives (sees, hears, etc.) an F directly, there is an F that one perceives directly.

2. Sometimes one perceives an F directly without there being any physical object that is F and that one directly perceives.

3. Sometimes one directly perceives non-physical objects (from 1-2).

4. The objects we perceive directly are always of the same kind (either physical or non-physical).

   Therefore, the objects that we perceive directly are always non-physical objects.

Claims along the lines of premise S1 have been treated as self-evident by a number of proponents of such arguments (c.f., Broad 1952, Moore 1910; 1905, Price 1932).

It is a version of what Crane (2005) and Robinson (1994: 32) call the "phenomenal principle". Premise S2 asserts the possibility of a certain kind of illusion. It is typically motivated by the observation that statements such as the following are sometimes true even though there is nothing physical in one's environment corresponding to what one is said to see:

- All he sees on his left is a grey blob. (Describing a neurological condition)
- I saw a pink elephant. (Describing the effect of a drug)
- She is seeing stars. (Describing the effect of a sudden head movement)

Premise S4 is sometimes described as the "spreading step". It warrants extending the conclusion drawn about illusory experiences to veridical experiences.

Harman's objection to argument S, as formulated his 1990 paper, turns in part on a distinction between two kinds of seeing or perceiving. Typically, when one says that one sees an F, this implies that there is an F. For example, if, looking at the parking lot outside my window, and confident that I am perceiving the world as it is, I sincerely declare that I see a green Volvo outside, my claim arguably entails that there is a Volvo I am suitably related to. Following Anscombe (1965), we can call this the *material reading* of the ascription. There is also a sense in which I can see a Volvo (or Volvo-like object) without there being any. This is the *intensional reading* of the ascription.[2] This is the kind of reading on which it is true that some people have seen pink elephants, stars, and gray blobs due to neurological dysfunction. One of Harman's key points in response to argument S is that premise S1 is not true if "perceives" is given an intensional reading, while premise S2 is only true on an intensional reading. He uses "see*" for the material reading and "see†" for the intensional reading:

> The argument from illusion starts from a case in which Eloise "sees" something brown and green before her, although there is nothing brown

[2] Anscombe writes "intentional", but I find "intensional" more appropriate.

> and green before her in the external physical world. From this, the argument infers that the brown and green she sees must be internal and mental. Now, if "see" is "see†" here, this is the fallacy already noted, like that of concluding that Ponce de Leon was searching for something mental from the fact that there is no Fountain of Youth in the external world. On the other hand, if "see" is "see*" here, then the premise of the argument simply begs the question. No reason at all has so far been given for the claim that Eloise sees* something brown and green in this case. It is true that her perceptual experience presents her as visually presented with something brown and green; but that is to say merely that she sees† something brown and green, not that she sees* anything at all. (Harman 1990: 36-7)

Harman's response might seem to turn exclusively on the material / intensional distinction, but I think it is intended to cut deeper than this. The material / intensional distinction is sufficient to address the particular form of argument at hand, but another point Harman makes in the last sentence extends his response to all arguments from illusion: the only sense in which Eloise can "see" something brown and green when nothing is in her environment is equivalent to saying that she experiences something brown and green. Experience, he further suggests, is intentional just like searches (so an experience is at best a state in which one sees†, not sees*). This claim puts Harman in a position to block not only the present argument but any parallel argument, whether cast in terms of "seeing" or not. In all such arguments, there will be premises to the effect that one cannot experience an F without there being one. According to Harman, this is not true because experience is intentional.

Here is how Harman's general response applies to argument S in particular:

Argument H: Harman's counterargument

1. Premise S2 is only true in the sense of "perceives an F" which picks out a phenomenal state.

2. Phenomenal states are intentional.

3. If (1) and (2), then it is possible to perceive an F without there being any F (in the sense of "perceive an F" which makes S2 true).

4. Premise S1 is only true in a sense of "perceive an F" on which it is impossible to perceive an F without there being an F.

   Therefore, premise S2 is only true in a sense of "perceives an F" on which premise S1 is false.

This statement of Harman's reply brings into sharp relief the role of representationalism: its role is to supply premises H2 and H3. H2 merely states that representationalism is true (or that a relatively weak kind of representationalism is true), so the test for a formulation of representationalism is whether it supports H3 to satisfaction. Is this true of virtualism?

Suppose that "perceiving an F" picks out a phenomenal state as H1 requires, and suppose that phenomenal states are as virtualism says they are (i.e. assume the antecedents of H3). In this case it would seem that "perceiving an F" picks out a state which obtains when one stands in R to a proposition involving F-ness. Since R is virtual, this would strongly suggest that one can perceive an F without there being an F. Virtualism therefore seems to supply premise H3. It does not obviously entail premise H3, but it supports it as well as any other representational theory. In the next chapter I will give a detailed account of intensional readings of perceptual ascriptions which establishes a more direct connection between virtualism and H3. For now, I think we can at least see that there is no need to say that experience is representational or intentional to block the argument from illusion: we only need

to say that it is virtual. Adding that relation R is a representational or intentional relation would not make the reply go any better.

Blocking the argument from illusion for sense data is part of what a theory must do in order to solve the problem of perception, but it is not everything. To solve the problem, a theory must account for illusion while preserving the immediacy datum. We must now ask: a) is virtualism compatible with the immediacy datum?; b) does virtualism really account for illusion?

One might think that if we experience propositions it follows that we don't experience external facts or states of affairs directly. Is not the veil of perception a veil of abstract objects on the virtualist view? But recall that we count states of affairs as propositions. I am using "proposition" as an umbrella term for all proposition-like things. On the virtualist view, what you are directly aware of (what you experience) when you perceive an external scene can perfectly well be the scene itself, not a surrogate: the scene is a state of affairs, which is a kind of proposition. The content of your experience can be the very same external scene that you perceive. If this state of affairs happens to be a fact, what you experience is a fact, and you experience it directly (I discuss this topic a greater length in section 7.2).

Of course, it is part of the theory that you could have the very same kind of experience—stand in R to the very same state of affairs—without the state of affairs you experience obtaining. But the point of virtualism is that this does not imply that what you experience in the veridical case is something else than the state of affairs that is before you. What this implies is that in experience you can stand in a relation to a state of affairs which does not obtain, i.e. that the experiencing relation is virtual. I think it is safe to say that virtualism at least preserves the immediacy datum and avoids the "veil of perception" consequences which make the sense-datum theory clearly unacceptable. On the virtualist view, you do not have to infer the existence of a three-dimensional world from a two-dimensional array of pixels or any other mimicry of the world. You simply believe the content of your

experience, and this gives you a belief about a three-dimensional world. Whether there is a kind of immediacy virtualism does not allow beyond this—and whether perception is immediate in this stronger sense—is a more complex question I will turn to in chapter 7.

Question (b) is also bound to elicit negative answers. In what sense does virtualism account for illusion? It might seem to provide less of an explanation than the theory that experience is representational or intentional, because this theory might seem to at least provide a model for understanding what goes on in illusion. But virtualism does that too. The model it provides is the general class of virtual relations. What virtualism loses in specificity, it gains in precision and clarity. Rather than saying vaguely that experience is like belief, it says exactly in what relevant respects it is like belief. Overall, virtualism illuminates the nature of illusion as much if not more than other general, ecumenical representational theories of consciousness.

One might say that conventional representational theories formulated in terms of intentionality or representation at least have the advantage of paving the way for further explanation by assimilating experiences to states of a kind there are many candidate reductive explanations for. But we have found no suitable meaning for "intentionality" and "representation". A consequence of this is that it is unclear whether the explanatory target of "reductive theories of intentionality" really is the kind of intentionality or representation which some representationalists want to assimilate experiences to. In fact, some proponents of “theories of intentionality" (e.g. Dennett 1971 and Millikan 1984) make explicit that they are not trying to explain a pre-theoretically grasped phenomenon of "intentionality" but merely to characterize an interesting kind of state in terms of which behavior can be explained.

Setting this aside, a proponent of virtualism also has access to reductive theories of intentionality to explain the virtual character of experience: just like a proponent of conventional representationalism can claim that the intentionality of experience

is suitably explained by such and such a theory of intentionality, a proponent of virtualism can claim that the virtual character of experience is suitably explained by the fact that the virtual relation constitutive of experience is the kind of relation labeled "intentionality" or "representation" in such and such a theory of "intentionality" or "representation".

2.2 The transparency of experience

Aside from accounting for illusion, another role representationalism is generally agreed to play is that of explaining the transparency or diaphanousness of experience.[3] A number of different transparency theses have been discussed in the literature. I will once again begin with what Harman says in his seminal article. This oft-quoted passage provides a natural starting point:

> 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. (Harman 1990: 39).

Harman is making two claims: one negative and one positive. The negative claim is that, if one tries to attend to intrinsic features of a visual experience, one will fail. The positive claim is that, if one tries to attend to intrinsic features of a visual experience, one's attention will be drawn to what we can on a first gloss describe as features of external objects. I will discuss each point in turn.

2.2.1 Negative transparency

When Harman talks about intrinsic features, he does not just mean intrinsic properties. He is not saying, for instance, that introspection fails to show experiences'

[3]Transparency is sometimes taken to be a key premise of some deductive arguments for representationalism (c.f. Stoljar 2004), but it is also sometimes presented as an explanandum for this theory (c.f. Tye 1995).

durations or levels of complexity, even though these are arguably intrinsic properties of experiences. He also does not mean simply intrinsic qualitative properties, at least not in the common, broad sense of "qualitative" which is roughly equivalent to "not quantitative". Again, Harman is not trying to suggest that experiences cannot introspectively seem brief or complex, and these are arguably intrinsic qualitative properties. What he means by "intrinsic features" is monadic qualitative properties of the kind we experience in perception—the properties our perceptual experiences are typically said to be "of" (e.g., redness). I will refer to properties of this kind as "qualia".[4] Harman's negative point is that, as a general rule, introspection does not present phenomenal states as instantiating qualia. This is the negative transparency thesis representationalism is often taken to shed light on.

Note also that representationalism is not supposed to explain negative transparency by showing that introspection *misleadingly* presents qualia as belonging to external objects instead of experiences. It is supposed to explain transparency by explaining, predicting, or otherwise shedding light on the fact that experiences do not instantiate qualia at all.

Now, explaining the non-existence of something is different from explaining the existence of something. One does not normally explain why a property or a state of affairs does not obtain by uncovering a causal mechanism or postulating principles in virtue of which it does not obtain. As a general rule, explanations of negative facts (to the extent that we can speak of explanations at all) take the form of accounts of reality which remove any reason for positing that the relevant entities obtain. For example, modern physics in a sense explains the fact that there is no phlogiston by providing an account of thermal phenomena which removes the need to posit phlogiston. Representationalism can at best explain the absence of qualia in this kind of way, because the view is universally conceived of as a thesis about what experience is, not what it is not. Whatever the representationalist says experience is

[4]Qualia, in this sense, are not the same thing as Block's (1998) Qualia or Chalmers' (1996) qualia, or Dennett's (1988).

(e.g., a kind of informational process), it cannot be directly inferred from this that experiences do not also have intrinsic qualia: for all we know, experience could be an informational process, yet be accompanied by qualia.[5] (Indeed, Shoemaker (1994) essentially endorses this view.) What representationalism can do is remove the need to posit intrinsic qualia.

Why might one think that experiences instantiate qualia—colors and sounds, or similar properties (e.g. red')? A possible response is that one can verify that experiences have such properties introspectively. But this is simply the negation of negative transparency. Representationalism is supposed to shed light on negative transparency by undercutting the motivation for positing intrinsic qualia. If this motivation were the negation of negative transparency itself, the representationalist would be in a position where a) her explanation of T (negative transparency) is that N (no intrinsic qualia) is true and implies T; b) her explanation of N is that not-T is false (T is true). In other words, representationalists' explanation of negative transparency, insofar as they have one, would be negative transparency itself. This cannot be how representationalism sheds light on negative transparency, so the motivation for positing intrinsic qualia which representationalism undercuts must be something else than that they are manifest in introspection.

The only other potentially historically influential motivation for intrinsic qualia I am aware of is the argument from illusion (and its variations). This is roughly the motivation outlined by Crane (2000) in his history of the concept of a quale. Crane essentially argues that the qualia theory is a version of the sense-datum theory according to which the bearers of the properties we experience (sense data) are our experiences themselves. His discussion also suggests that the theory's motivation is the argument for sense data combined with a rejection of mental particulars distinct from familiar mental events. Harman (1990) also suggests that the argument from illusion is the primary motivation for intrinsic qualia. If this reconstruction of the

[5]This point is discussed at greater length in Seager and Bourget (2007).

motivation for intrinsic qualia is correct, virtualism's ability to block all arguments along the lines of argument S should already ensure that it undercuts the motivation for intrinsic qualia (hence explains negative transparency).

2.2.2 Positive transparency

Let us now turn to positive transparency. The claim, as put by Harman, is something like this: as a general rule at least, if you try to attend to intrinsic features of your experience, you will end up attending to what might be described as features of external objects. This gloss of positive transparency is problematic because it is unclear what is meant by "features of external objects". I can think of four reasonably likely interpretations of this phrase:

1. Features which are as a matter of fact instantiated by external objects
2. Features which seem to you to be features of external objects
3. Features which you experience as features of external objects
4. Features of the external kind

The "seem" in option 2 should be read epistemically (see p. 32 for the distinction between epistemic and non-epistemic readings of "seem"). It is only on this reading that it differs significantly from option 3. Option 4 can also be given multiple interpretations. For the sake of concreteness, we can say that a property is external (is of the external kind) just in case a) its instantiation requires the instantiation of spatial properties; b) it is publicly observable if observable at all; c) it is mind-independent in the sense that its instantiation does not require the existence of a mind. These three conditions seem to me to capture a widespread notion of external property, but the details are not very important for our purposes.

Option 1 does not seem acceptable. Read this way, positive transparency would imply that the features we attend to in introspection are always actual features of

external objects. This would imply that positive transparency is not always true of hallucinatory and illusory experiences, an implication which representationalists seem to reject (c.f. Tye 2002a: 45-51).

It is worth looking at alternative statements of positive transparency to see if any of the remaining options emerges as a canonical interpretation.

> In turning one's mind inward to attend to the experience, one seems to end up concentrating on what is outside again, on external features or properties. (Tye, 1995: 136)

Assuming option 1 is not charitable enough, Tye's statement is most naturally understood along the lines of 2 or 4. By contrast, this statement by Chalmers is most naturally interpreted as an instance of either 3 or 4:

> [...] the central datum of transparency is that when we attempt to introspect the qualities of our experiences (e.g. phenomenal redness), we do so by attending to the qualities of external objects (e.g. redness) (Chalmers, 2004)

Martin's gloss of transparency suggests 2 or 4 (4 with a more permissive definition of external properties as mind-independent properties):

> At heart, the concern is that introspection of one's perceptual experience reveals only the mind-independent objects, qualities and relations that one learns about through perception. (Martin, 2002)

It is not clear that there is a canonical interpretation of positive transparency, so we will have to see if we can explain all three available interpretations. Let us start with interpretation 4, which appears to be closest to consensual. On this interpretation, the positive transparency observation is that if one introspects one's perceptual experiences, one will end up attending to properties of the external kind. That is what I will mean by "positive transparency" until further notice.

It should come as no surprise that neither virtualism nor any plausible general representational theory explains positive transparency in the strong sense of entailing it or allowing us to deduce it. Representationalism is supposed to be a view about all experiences, while positive transparency is a claim about specifically perceptual experiences and has never been construed otherwise. For example, nobody thinks that positive transparency is true of headache experiences, cognitive experiences, and emotional feelings (as far as I know). Positive transparency highlights a peculiarity of perceptual experience, so we should not expect a general account of consciousness to explain it. There is however a significant sense in which virtualism explains positive transparency: from virtualism and auxiliary assumptions which are plausible in light of virtualism, we can infer positive transparency. The relevant auxiliary assumptions are the following:

The external content thesis: The contents of perceptual experiences involve external properties.

Closure under composition of awareness (CCA): Awareness of a relational state such as standing in R to α requires awareness of the components of α.

The external content thesis is more or less forced on us by virtualism. For if virtualism is true, nothing but its content can distinguish a perceptual experience from a non-perceptual experience.[6] The only contents which are plausibly characteristic of perceptual experiences are those involving such properties as colors and shapes understood as mind-independent properties—straightforwardly external properties. This means that the contents of perceptual phenomenal states involve external properties.

The rationale behind CCA is that nothing distinguishes two relational states which involve the same relation except for the relata they involve. Take for example these two states:

[6] "Perceptual" here indicates a certain kind of phenomenology, not a certain kind of etiology; there can be hallucinatory perceptual experiences.

owning a Volvo

owning a Mazda

There is no difference between these states aside from the object owned. It therefore seems that, to be aware of or attending to Bob's state of owning a Mazda (as opposed to his state of owning something), one would have to be aware of or attending to a Volvo. This reasoning can be repeated at the level of the relata when these are states of affairs. Take for example these two propositions:

There is green jelly on the table.

There is blue jelly on the table.

There is no difference between these two propositions aside from the colors they involve. It therefore seems that awareness of one of these states of affairs in particular requires awareness of the color involved in it—some kind of cognizing or consciousness of the color, depending on how "awareness" is understood exactly. Assuming this is right, being aware of standing in R to the proposition that there is green jelly on the table would require being aware of greenness.

Positive transparency falls directly out of virtualism given CCA and the external content thesis: from the external content thesis and virtualism, we know that every perceptual phenomenal state is a relational state of standing in R to a proposition involving external properties; from this and CCA, it follows that awareness of a perceptual phenomenal state requires awareness of external properties.

This explanation of type-4 transparency extends to other kinds of positive transparency. Take first the claim that when one introspects a perceptual experience, one ends up attending to properties which one experiences as properties of external objects (interpretation 3). Presumably, to experience a property as a property of an external object is to experience it as a property had by something which has properties which make the object external. To a first approximation, it seems plausible that these external-making properties are precisely properties of the external kind.

If this right, attending to an external property that one experiences is ipso facto attending to a property that one experiences as a property of an external object. That is to say that type-4 transparency implies type-3 transparency.

This is not true of type-2 transparency (the claim that if one introspects a perceptual experience, one will end up attending to properties which will seem to one to be properties of external objects). Personally, I have doubts about type-2 transparency. It seems to me perfectly possible to suspend all belief about the external world while introspecting one's experiences, and in these circumstances it would not epistemically seem to one that the properties one finds in introspection are properties of external objects. Still, we can explain the observation to the extent that it is correct based on the fact that the world normally seems to us epistemically to be the way it is represented in our experience. Given the external content thesis, CCA, and this additional assumption, we can derive type-2 transparency for the normal case. Virtualists are not alone in having to posit that the world normally seems to us to be the way we experience it in order to explain type-2 transparency, because this posit is not part of the idea that experiences are intentional or representational (there are representations which are more desire-like than belief-like).

2.3 The place of consciousness in nature

One of the most important roles representationalism can play is to guide our inquiry into the relation between consciousness and the physical world. In this section I ask whether virtualism can play this role as well as other general, ecumenical representational theories. I begin by distinguishing two questions regarding the place of consciousness in nature.

2.3.1 The metaphysical and scientific questions

There are two central questions regarding the relation between phenomenal and physical states:

The metaphysical question What metaphysical relations hold uniformly between phenomenal and physical states?

The scientific question Which phenomenal states accompany which physical states?

The correct answer to the metaphysical question will tell us which of identity, supervenience, causation, realization, determination, ecetera hold between phenomenal states and physical states. The correct, complete answer to the scientific question may be expected to specify a systematic mapping between phenomenal states and physical or functional states. This mapping would determine which phenomenal states go with which physical states. As Chalmers (1996: 214) points out, a truly satisfactory answer to this question would arguably have to fit on a t-shirt.

The metaphysical and scientific questions are clearly distinct. A number of philosophers appear to think that they have solved the metaphysical puzzle without addressing the scientific question. They think, for example, that considerations pertaining to causal closure show that phenomenal states are identical to physical states, but they offer no account of which phenomenal states are identical with which physical states. Conversely, neuroscientists do not as a general rule seem to be preoccupied by the metaphysical question. For them it makes little difference whether phenomenal states are identical to physical states or not. Their primary concern is to discover which phenomenal states go with which physical states, because that is all that one needs to know in order to predict all relevant conscious events.

While the metaphysical and scientific questions are logically distinct, it is not unreasonable to expect an answer to one to bear on the other. In particular, it seems likely that a fully satisfactory answer to the scientific question—one that relates

phenomenal states and physical states via an equation that fits on a t-shirt—would militate for the identity theory of phenomenal and physical states. For one thing, it would alleviate worries about multiple realizability. It would also attenuate the impression that the mind-body relation is arbitrary, which speaks against identity. For example, the theory could conceivably be taken to suggest that inverted qualia cases are not metaphysically possible. It could potentially do this by giving us a better grasp of the nature of consciousness and exhibiting its structure. The theory could not show with certainty that Chalmers' (1996) zombies are impossible (not if it is formulated as a nomic generalization, as it should be) but it would plausibly sway a number of skeptics.

I will now survey the applications of representationalism with respect to the scientific and metaphysical problems, and ask whether virtualism lends itself to the same applications.

2.3.2 Virtualism and the scientific problem

One of the principal attractions of representationalism is its ability to give us a better handle on the internal structure of consciousness. Virtualism does this at least as well as any form of representationalism. First, it ascribes a relational structure to phenomenal states. This alone is potentially illuminating. As we will see in chapter 5, this simple fact can shed light on puzzles about the nature of sensory modalities and the relation between perceptual and non-perceptual experiences. The fact that the relata of R which characterize phenomenal states according to virtualism are propositions is also promising. Propositions exhibit internal structure and can be described formally. This is what helps with the regimentation problem mentioned in introduction (which is part of the scientific problem). Virtualism provides a framework in which to articulate and regiment the natures of phenomenal states. This is just what we need in order to formulate clear, precise, and fully general hypotheses about how phenomenal states relate to each other and physical states. Here

are some examples of interesting generalizations which can only be expressed in a framework like that provided by virtualism (R is the relation posited by virtualism):

Phenomenal consistency It is impossible to stand in R to inconsistent propositions.

Phenomenal composition Necessarily, if *x* stands in R to both P and Q, *x* stands in R to P∧Q.

Phenomenal decomposition Necessarily, if *x* stands in R to P∧Q, *x* stands in R to both P and Q.

I don't know that any of these principles is true, but I find them prima facie plausible.[7] In any case, the important is that they illustrate the potential explanatory power of the framework. Phenomenal composition and consistency, for example, together entail that if one stands in R to P (experiences P), one cannot experience ¬P. This might provide a fundamental explanation of the gestalt shifts one experiences with Necker cubes and the like (this would explain why we never see the cube both ways at once).

Another potential application I see for these principles is an explanation of the unity of consciousness. I take it (following Bayne and Chalmers 2003) that it is necessary and sufficient for one's phenomenal states to be unified at a time that they all be entailed by some phenomenal state one is in at that time (the total phenomenal state one instantiates at that time). This by itself does not explain unity but merely defines it more precisely.[8] To explain the facts about unity so defined, we must at least explain 1) why all of one's phenomenal states are always entailed by

[7]There has been some discussions of putatively inconsistent experiences, e.g. in Crane (1988), O'Shaughnessy (1957), and Sorenson (2002). I find that it is generally possible to explain seemingly inconsistent experiences as rapid alternations between distinct, internally consistent experiences. This interpretation is strongly suggested by the fact that the consistency of experience is as a general rule extremely robust: no matter how much drug one takes or how tired one is, somehow one always ends up with a logically consistent image of the world. This is even true in dreams, which are otherwise quite unconstrained. What one experiences might not always square with one's conception of the world, but consciousness is remarkably reluctant to present us with an impossible world.

[8]Of course, one might want to recognize other kinds of unity.

some phenomenal state one is in and 2) why tuples of people or other aggregates of conscious beings do not have experiences which unify the experiences of their members, e.g. why the mereological sum of me and my mother (call it MAM) does not have phenomenal states unifying our experiences. That is to say that we must explain both the unity there is and the unity there is not.

Observation 1 is straightforwardly explained by phenomenal composition and decomposition together: by composition, someone who stands in R to both P and Q also stands in R to P∧Q; by decomposition, this state itself entails standing in R to P and standing in R to Q. This explanation is also consistent with observation 2, because *my* component phenomenal states are not the phenomenal states of MAM, so it does not follow by composition that MAM has any experience which unifies them.[9]

A related application of the virtualist framework would be to explain the productivity of phenomenal consciousness. Like thought, experience comes in a huge variety. It seems plausible that the variety of experience is generated through some productive combination mechanism. One can easily imagine a parallel to Fodor's Language of Thought strategy applied to phenomenal states within the virtualist framework. We can imagine that there are basic ingredients of phenomenal states

[9]Of course, this explanation seems superficial because it is glossed in terms of persons. This makes it a little too easy to give an explanation consistent with observation 2. But just say that the following is true:

Phenomenal contagion If a part of a physical system S stands in R to P and no part of S stands in R to a proposition incompatible with P, then S stands in R to P.

The idea behind phenomenal contagion is that phenomenal states spread to higher structures as much as they can without risk of generating inconsistent states. Assuming that phenomenal states are had at the most basic level by parts of persons, contagion, composition and decomposition together predict roughly what we find in introspection: by contagion, and given that all the experiences had by parts of me are consistent, I have all these experiences; by composition and decomposition, they are all unified. What is remarkable is that this simple and elegant explanation of observation 1 is in line with observation 2: even though MAM and I share parts which instantiate phenomenal states, MAM cannot have experiences which unify my mother's and mine through contagion because our experiences are inconsistent. They are inconsistent because we always experience different, incompatible properties as being instantiated at locations which are subjectively the same for us. For example, my current visual content says that what is "right in front" has such and such quality. My mom's visual content ascribes different, incompatible properties to what is right in front of her. If both of our contents are about a generic "what is in front"—the same thing in both contents—, they are inconsistent.

associated with the various constituents of the propositions or states of affairs we can experience (e.g. properties), and that all phenomenal states are formed by combining these ingredients. The content of a complex phenomenal state would be determined by the contributions of its constituents in much the same way that the content of a sentence is determined by the semantic values of its constituents.

But enough speculation. My point is that it is not hard to see how the structure virtualism adds to consciousness could be put to use. The importance of this point cannot be overemphasized. Without a framework like that provided by virtualism, consciousness is shapeless and elusive. To make progress toward a general solution to the scientific problem, we need a language in which to state precise generalizations about consciousness. Virtualism promises this: it promises a formal framework in which to state the Newtonian laws of consciousness—perhaps not the principles stated above, but hopefully similarly simple principles.

Some reductive representational theories purport to do more for the scientific problem than virtualism. Dretske's (1995) and Tye's (1995) theories, in particular, are potential solutions to this problem, because they provide general mappings between phenomenal states and physical-functional states. But there are good reasons to reject these theories, some of which were summarized in the preceding chapter (section 1.2). If they are false, they do nothing for the scientific problem. Our concern, in any case, is to assess whether virtualism can help with the scientific problem as much as the other general representational theories most representationalists can accept.

2.3.3 Virtualism and the metaphysical question

Before we ask whether virtualism can do the job of representationalism with respect to the metaphysical problem, we must first clarify what this job is.

There are two main ways representationalism could conceivably help with the metaphysical problem independently of what it can do for the scientific problem.

First, it could provide independent support for one metaphysical theory or the other. For example, it could help us see that the identity theory is true (or false). Representationalism is more often than not taken to support physicalism, not dualism, so I will only consider it in this role.[10] Given virtualism's complete agnosticism regarding the realizers of phenomenal states, it should also be fairly clear that it is just as good as any other form of representationalism as far as supporting dualism goes. Second, representationalism could undercut the evidence against one metaphysical view or the other. This is different from supporting one view or the other: sometimes showing that an argument for P is unsound does not provide any independent evidence for ¬P but merely undercuts the evidence for P. Here I am not aware of purported applications which are friendly to dualism. What has been suggested is that representationalism helps undercut the zombie and knowledge arguments against physicalism (c.f. Tye 1999, Jackson 2003). Representationalism can also help with the metaphysical problem in virtue of helping with the scientific problem, but we already discussed the relation between virtualism and the scientific problem.

I don't think representationalism provides much support for physicalism. Of course, physicalist versions of the view such as Dretske's, Tye's, and Lycan's theories directly imply physicalism.[11] Aside from the fact that there are good reasons to reject these theories which do not rely on general arguments against physicalism, there is the problem that they do not provide independent support for physicalism: no one who finds physicalism or functionalism unattractive will find them attractive. If these theories do not help convince the skeptics, they do not provide support for physicalism in any interesting sense. It is also true of nonreductive versions of representationalism (including virtualism) that they provide no support for physicalism independently of whatever impact they might have on the scientific problem. One might have thought that identifying phenomenal states with intentional states

[10] An exception is Pautz (2006), who argues for dualism from an intentionalist perspective.

[11] To keep things simple, I count functionalism as a kind of physicalism.

would help. But no one thinks that phenomenal states are mere intentional states (except maybe for people who think that intentional states are all phenomenal, but this would defeat the reductionist agenda). Extant nonreductive representationalist views claim at best that phenomenal states are a special phenomenal kind of intentional state. Anyone who sees a problem with phenomenal states being physical or functional should remain unmoved.

Whether representationalism undercuts arguments against physicalism is a more complicated matter. In practice, the evidence for dualism reduces to a reasoning of this form:

The generic case for dualism

1. Phenomenal states intuitively seem to be X.

2. That phenomenal states intuitively seem to be X is good evidence against physicalism.

   Therefore, there is good evidence against physicalism.

This is not how the arguments for dualism are normally presented, but what they boil down to in practice—the points we have ended up debating. Officially, the main arguments for dualism have this form: "phenomenal states are X; if phenomenal states are X, physicalism is false; therefore, physicalism is false". But the debate has taken a psychological turn due to the nature of the relevant Xs. As it is, the key question is whether how phenomenal states intuitively seem to be is a good indicator that dualism is the correct view.

In the case of Chalmers' zombie argument, the first premise of the case for dualism is that zombies intuitively seem conceivable (X is the property of not being a priori entailed by physical states). In the case of Jackson's knowledge argument, the first premise is that it appears possible to know all the physical facts without knowing certain facts about the nature of consciousness (X is the property of having a nature that is not known automatically upon knowing all the physical facts).

Both starting points are highly plausible at least on some reading. In each case, a favorable reading of the first claim can be adopted and all the burden of the argument shifted onto the second claim, which asserts that the intuitively plausible fact indicates the falsity of physicalism.

The main reasons for disagreeing with the second claim which have been advanced fall into three categories: appeals to broadly speaking semantic features of the relevant statements or concepts (e.g., a posteriori necessity or indexicality; Perry 2001); appeals to a special kind of concept (phenomenal concepts; c.f. Loar 1990, Lycan 2003; 1996a, Tye 2000: 30); appeals to the hypothesis that dualist intuitions have their source in a systematic conflation of experiences with their objects (Jackson 2003; to some extent Harman 1990). The first strategy has not to my knowledge been associated with representationalism. The second strategy is endorsed by representationalists Lycan and Tye, but it does not obviously draw on representationalism. (The strategy is to claim that the concepts we use in grasping experiences are "recognitional" or relevantly like perceptually grounded concepts.) The last approach, however, is quintessentially representational.

In a nutshell, the claim is that we systematically conflate the properties we represent in experience with properties of our experiences. Since, the argument goes, the properties we represent, at least in the case of color experience, are not entailed by physical facts (they are not instantiated at all), this gives rise to the illusion that phenomenal facts are not entailed by or identical to physical facts, illusion which manifests itself in the intuitive appeal of Chalmers' and Jackson's starting points. Here is Jackson on Jackson:

> There is a redness about sensing red (a yellowness about sensing yellow, and so on). We naturally think of the redness as a property we are acquainted with when we sense red and as the property Mary finds out about on her release. We may want to distinguish redness as a property of objects from redness as a property of an area of our visual field,

> perhaps using 'red*' for the latter. Either way, what it is like is, on the picture, a matter of having redness or redness*, knowing what it is like is knowing about redness or redness*, and the knowledge argument is an argument to the conclusion that Mary does not know about redness or redness*—that is, about the property we are, according to the picture, acquainted with when we sense red. [...] intentionalism tells us that there is no such property. To suppose otherwise is to mistake an intentional property for an instantiated one. (Jackson 2003)

Jackson goes on to argue that redness or redness*, the property that dualists such as his old self conflate with a property of experiences, is not even instantiated in this world. This property is not entailed by physical facts, but that is compatible with physicalism because it is a mere intentional property.

The negative transparency thesis does most of the explanatory work for Jackson. Indeed, the above passage follows his discussion of negative transparency. In a nutshell, Jackson's reply to his old self is that a) negative transparency is true; b) negative transparency shows that experiences do not instantiate redness or redness*; c) the knowledge argument is about redness or redness*; therefore, the knowledge argument does not bear on consciousness. If this is how representationalism undercuts the evidence for dualism, any representationalist view which supports negative transparency can do the job (including virtualism).

I must say that I do not find Jackson's response entirely satisfactory. It does not seem hard to keep the distinction between instantiated properties and represented properties straight, once one has seen it. Suppose we think of experiences of red as phenomenal representations of redness which do not instantiate any redness-like properties (in some sense or other of "representation"). Doesn't it still seem that experiences of red fail to be entailed by physical states, or that one could know all the physical facts while not knowing what it's like to be in such a state? While the representationalist view can provide an explanation of one possible kind of cogni-

tive illusion, it is not clear that we cannot avoid this illusion while preserving the intuition that fuels dualist sentiments.[12]

There is however a further move one can make within the representationalist framework beyond Jackson's reply. This one involves positive transparency. If positive transparency is true, we systematically attend to external features such as redness when attending to our experiences. That is to say that we systematically attend to features which are not parts of our experiences but merely represented by them. This observation could go some way toward explaining our purported tendency to conflate our phenomenal states with properties like redness which are neither physical nor instantiated. It might, after all, be very hard if not impossible to keep the distinction between phenomenal states and their objects straight. This response would be further reinforced if it turned out that grasping an experience in order to reason competently about it required attending to it introspectively or attending to a "faint copy" in imagination (Tye 2000, ch.2 suggests something like this).

The foregoing strikes me as a potent reply to arguments against physicalism, and it works just as well on the virtualist view as on any other representationalist view. The reply is highly speculative, but it seems to be the best representationalism enables (maybe not the best available to physicalists). I conclude that virtualism, with plausible auxiliary assumptions, can help physicalism at least to the extent that representationalism may reasonably be taken to help it.

2.4 The place of consciousness in the mind

A theory which assimilates experience to a species of "intentionality" might naturally be expected to shed light on the relation between consciousness and propositional attitudes. There are a number of central questions one could expect such a

[12] Alter (2006) presents an objection along these lines.

theory to throw new light on, including the following two:

The epistemic role of experience What role do experiences play in the justification of beliefs?

The semantic role of experience What role do experiences play in grounding the contents of propositional attitudes?

Virtualism might seem less promising than other representational theories as far as answering these questions goes because it does not explicitly ascribe a common nature to experiences and propositional attitudes. But here it is important to keep in mind that other theories fail at this too, for reasons covered in chapter 1. Unless experiences are implausibly said to be canonical attitudes such as beliefs and desires, there is no clear content to the claim that they are intentional states or attitudes.

Of course, that is only to say that virtualism does no worse than other representational theories. The virtualist can in fact do better. I won't be able to explain how in detail here, but I want to sketch the overall picture.

It is best to start with the semantic role of experience. Most theorists would agree that experience is not only a source of evidence about the world, but also a source of content: it is largely through experience that we acquire the ability to think about the world that surrounds us. Although this was not the point of it, this fact is brought into sharp focus by Jackson's Mary thought experiment: there is a strong intuition that Mary cannot acquire the concept RED—the normal one that normally sighted people have—until she experiences red. The problem for the virtualist is to explain why this might be.[13]

I suggest that experience furnishes the mind with content by realizing canonical propositional attitudes. We have seen that experiences are not canonical propositional attitudes, but this does not mean that they cannot be realizers or constituents

[13]There is also another problem, that of grounding reference to particulars. This will be discussed (also briefly) in chapter 8.

of canonical propositional attitudes. Suppose for example that you have an experience of a certain colored shape in front of you. This experience by itself does not constitute a belief. However, it seems to me that it could well be an important player in the causal network of states (occurrent and potential) which constitutes or realizes your belief that there is something colored in front of you. On the output side, at least, it can have all the (defeasible) inferential connections we would expect a belief-realizing state to bear to other states.

We get a more rounded picture of the role experiences can play when we consider non-perceptual experiences. Mental imagery, whether visual, auditory or proprioceptive, is extremely pervasive. On the virtualist view, these "imagistic" experiences are of a kind with perceptual experiences. As I will argue in chapter 6, plausibly the only essential difference between them and perceptual experiences is that they have more general contents (for example, they represent determinables such as redness rather than specific shapes of red). If we grant this, we can also grant a more important role to experience in realizing propositional attitudes: imagistic experiences can provide the contents of the general beliefs and desires which they realize through suitable connections with other states.

Cognitive feelings also seem to play a major role in cognition. I am thinking of the feeling of deja vu, the feeling that something is right (or not right), the feeling of being confused, the feeling of similarity, the feeling of wanting something, etc. These feelings are like signposts we follow as we reason our way through the problems of life. For example, if you ask me, "are you sure about this?", I will look inside for that feeling of conviction. I know whether I have met someone before by the feeling I have while looking at them—do they feel familiar or not?[14]

We can distinguish two kinds of propositional attitude which derive from phenomenal states. The first, which we could call *conscious attitudes*, are attitudes

[14]Goldman (1993) argues along these lines.

which consist in having suitably related occurrent and potential experiences.[15] Having an experience of red suitably related to a feeling of conviction, sensory imagery, and all other phenomenal states one is in a position to token might be an example of a conscious attitude. Conscious attitudes can themselves realize other attitudes through their interconnections and relations they stand in to non-experiences (e.g. entities in one's environment). These other attitudes could be called *external attitudes* (they are external to consciousness). A possible example of an external attitude would be a thought that H_2O is F which one has in virtue of a) having a (conscious) thought to the effect that the watery stuff around here is F; b) the fact that the watery stuff around here is H_2O; c) further dispositions which amount to one's being interested primarily in the basic structure of the watery stuff around here as far as F-ness goes. I talk about how certain intentional states can derive from others in more detail in Bourget (forthcoming).

When I say that experience furnishes the mind with content by realizing attitudes, I don't mean to imply that thinking requires the capacity for phenomenal consciousness. This is a further claim I am not making here (though I am sympathetic to it). What I am suggesting is that our propositional attitudes are in actual fact largely realized by phenomenal states. It seems to me that this position is a nearly inevitable extension of the widely endorsed functionalist and dispositionalist views of attitudes.[16] If you hold such a view, whether you are an internalist or an externalist about content (but especially if you fall in the former category), you should agree that experience is a central realizer of thought: there is so much of it, it cannot but play a central role in shaping the high-level causal patterns which make up thought.

[15]Conscious attitudes should not be conflated with what I call "phenomenal attitudes" in Bourget (forthcoming). Phenomenal attitudes are phenomenal states which are also propositional attitudes. I don't think there are phenomenal attitudes (at least no canonical propositional attitudes which are phenomenal states), but I have no qualms with conscious attitudes.

[16]Schwitzgebel (2002) makes a good case for a dispositionalism about attitudes which gives a central role to phenomenal states.

An extended discussion of the epistemic role of experience would take us too far away from our topic, but it is not hard to see in outline which way a virtualist should go on this matter. The key concern is that, as Davidson famously puts its, only a belief can justify a belief. Since phenomenal states are not more belief-like than desire-like on the virtualist view, it might seem that they cannot justify beliefs directly. They could still provide evidence for beliefs in the same kind of way that footsteps provide evidence that someone has passed by, but this sort of picture of how experiences justify beliefs about the external world is widely held to be untenable.[17]

The solution to this problem again lies in functionalism about attitudes. While no single experience constitutes a belief on its own, groups of experiences, some actual and some potential, with their causal connections, may be said to either realize or be beliefs. Functionalism about attitudes offers an alternative to the two preceding accounts of the epistemic role of experience: in the basic cases of perceptual justification which are of interest here, individual experiences neither directly justify beliefs nor serve as evidence for beliefs; rather, they *constitute* beliefs when present in sufficient numbers and suitably connected. While this view does not have all the epistemological advantages of a direct justification theory, it does seem to avoid the main pitfalls of the traditional alternative.

I am aware that the preceding accounts of the semantic and epistemic roles of experience are all too sketchy. I offer them primarily as proofs of concept for more sophisticated accounts compatible with virtualism. I will return to the content-grounding role of experience in sections 6.5 and 8.2.

[17] See Brewer's (1999) extended case against this view, or Johnston (2006) on the "Wallpaper view".

2.5 Summary

In this chapter I have tried to show that virtualism can play the theoretical role representational theories should be expected to play: to account for the possibility of illusion without sense-data, to explain negative and positive transparency, to provide a framework in which the scientific problem can be tackled fruitfully, to offer a possible explanation of dualist intuitions, and to shed light on the interactions between the different departments of the mind.

Part II

A case for virtualism

Chapter 3

Experience in plain language

Having introduced virtualism and illustrated some of its applications, I will now make a case for the theory in this and the next two chapters. There is no shortage of objections and alternatives to representational theories of consciousness. I don't expect the argument I will present to obliterate all sources of opposition. However, I think this argument provides good independent evidence: it ought to have considerable weight when we assess the totality of evidence at the end of the day. I will turn to objections in chapters 6-8.

One major difficulty with consciousness which must be attended to at the outset is that the language we (philosophers) use to talk about it is not well grounded in everyday speech. The definition of experience I have been working with, in particular, is largely ostensive. We cannot use the traditional analytic and linguistic tools with terms which are not grounded in everyday speech. My aim in this chapter is to present a way of describing phenomenal states precisely in everyday language. I refer to the thesis to be defended in this chapter as *the perceptual conception of sensory experience* (PCSE). The perceptual conception of sensory experience is the starting point of my case for virtualism.

Overview of the case for virtualism

1. The perceptual conception of sensory experience is correct. (this chapter)

2. If the perceptual conception of sensory experience is correct, sensory virtualism is true. (chapter 4)

3. If sensory virtualism is true, virtualism is true. (chapter 5)

   Therefore, virtualism is true.

To a first approximation, the perceptual conception of sensory experience says that sensory phenomenal states can be referred to as episodes of seeing, hearing, tasting, ecetera, on intensional readings of these expressions. The perceptual conception of sensory experience is only a thesis about sensory experience, and I will be concerned exclusively with sensory experiences and phenomenal states in this and the next chapter. For the sake of legibility, I will use "phenomenal states" and "experiences" without qualification to refer to sensory phenomenal states and experiences. Sensory experiences include experiences in the familiar sensory modalities (touch, sight, etc.) as well as bodily experiences such as pain experiences. For present purposes I do not count sensory imagination as a kind of sensory experience.

3.1 Intensional perceivings

In chapter 2 we saw in passing that perceptual ascriptions can lend themselves to two kinds of reading: intensional and material readings. On a material reading of "α sees an F" (see_m), this statement can only be true if there is an F that α sees. By contrast, on the intensional reading (see_i) the statement can typically be true even if there is no F that α sees. It is worth considering more examples of intensional uses of perceptual verbs. To start with "see":

· I see a strange shape on my left. I think my retina is damaged.

- I see colored shapes spinning in front of me, but I know there are no such things there.
- I see flashes all over the place—will you stop poking my brain with this electrode!

Anscombe (1965) gives examples which illustrate intensional uses of "hear" and "smell":

- With this hearing aid, when you talk I hear some screeching noises; no low tones and the consonants are very indistinct.
- I hear a ringing in my ears.
- I keep on smelling the smell of burning rubber when, as I find out, there is no such thing.

For "feel" and "perceive":

- I feel something in my back even though there is nothing there.
- I feel a painful sensation in my leg, but I don't have a leg anymore!
- I perceive a blue expanse when I put the lenses on.

Intensional readings of perceptual verbs are available in other languages than English. In French, for example, direct translations of the previous statements are all perfectly intelligible and normal. To illustrate:

- Je vois une forme bizarre à ma gauche; je crois que ma rétine est endommagée.
- J'entends un scillement dans mes oreilles.
- Je sens une sensation douleureuse dans ma jambe, mais je n'ai plus de jambe!

One might think that "feel" is a special case in that there is no material reading of it which contrasts with the intensional reading. In the case of "see", everyone agrees that there is a reading of "I see an F" on which there has to be an F one is suitably related to in order for this statement to be true. In the case of "feel", there might seem to be no such reading, because it might seem that "I feel pain", for example, simply means that I am in pain, which never requires the presence of an object.

It is important to see that there is in fact a material reading of "feel". Recall the last time you hit a toe on a door frame. You most likely felt a painful sensation *in your toe*. If you had paid close attention to the phenomenology of your experience instead of trying to distract yourself from the pain by swearing and contorting yourself, you could have felt the sensation *spread* and *contract* with time. Most likely, you eventually noticed that the sensation you had been experiencing was *receding from your toe*. Throughout this experience, you could have referred to the sensation in your toe as a pain: you could have perfectly well said that you felt pain *in your toe*.

This example illustrates the fact that there are natural uses of "sensation" and "pain" on which these terms designate things of a sort that can be found in parts of the body and be characterized as having various spatio-temporal properties (e.g. contracting and spreading). "Sensation" and "pain", understood this way, do not refer to experiences or mental states: I can say that I have a sensation or pain in my toe in this sense, but I cannot say that I have an experience in my toe. So there seems to be readings of "sensation" and "pain" on which they pick out other things than experiences. The things they pick out furthermore seem to be the objects of bodily experiences, that is, the things that we naturally claim to be feeling when undergoing such experiences ("I felt pain in my toe"; "I felt an unpleasant sensation in my chest"). "Pain" and "sensation" can also be used to refer to mental states, but these two uses should not be conflated. Here, I will use "sensation" and "pain"

exclusively to refer to the objects of bodily experiences, in as much as possible.[1]

"Feel" is subject to the same ambiguity as "pain", because one can say that one "feels pain" in either sense of "pain": one can feel the mental state of pain in (roughly) the sense of having it, and one can feel a painful sensation in a different sense of "feel". It is this second use of "feel" which lends itself to either material or intensional readings. On the material reading, "I feel pain" can only be true if there exists a painful sensation that I feel; not so on the intensional reading.

It is worth noting that while intensional readings are weaker than material readings in one respect (they have a kind of "existence neutrality"), they are sometimes stronger in another: to a first approximation, to say that one sees_{i} an F is to say something substantive about how the world appears to one, but to say that one sees_{m} an F is not to say anything substantive about how the world appears to one. For example, one can see_{m} a ship as a blip on a radar screen, as a big ship-like object, or as a light source in a fog. Seeing_{m} an F might require that an F appears one way or another, but there are virtually no constraints on how it can appear to one. By contrast, intensional readings of perceptual ascriptions of the form "α ϕ-s an F" have substantive implications regarding how the world appears to one.

What, exactly, is the difference between intensional and material readings of perceptual ascriptions?

We standardly distinguish three marks of intensional contexts: substitution resistance, existence neutrality, and non-specificity. A context resists substitution when changing a term in it for a co-extensive term can result in a change in the truth value of the embedding sentence. On the standard account, a context is existence-

[1] Brentano notes the tendency to conflate experiences and sensations:

> A further basis for this illusion [that experiences of bodily sensations do not "present" anything, in Brentano's sense] is the fact that the quality which precedes the feeling and the feeling itself do not have two distinct names. The physical phenomenon which appears along with the feeling of pain is also called pain. ... we say that we feel pain in the foot. This is an equivocation, such as, indeed, we often find when different things are closely related to one another. (Brentano 1874: 84)

neutral when existential quantifiers and singular terms occurring in it do not have their usual existential import. Non-specificity is best explained through an example. Take statement 3.1, an example from Forbes (2008).

(3.1) Oedipus is looking for a member of his family.

The specific reading of this statement is given by 3.2.

(3.2) $\exists x(familyMember(Oedipus,x) \wedge lookingFor(Oedipus,x))$

By contrast, the non-specific reading of 3.1 is one on which it does not entail 3.2. On this reading, Oedipus does not care who he finds so long as the person is a member of his family.

While fairly standard, the preceding account of intensionality is not without problems. In particular, Forbes (2006; 2008) gives several examples of contexts which appear to satisfy some but not all of these marks of intensionality. Coburn (1977) makes a parallel observation regarding the intensionality of perceptual verbs.

Given that the three "marks of intensionality" can potentially come apart, we would do best to regard them as characterizing potentially independent kinds of intensionality. Rather than talk about intensionality without qualification, we should talk about contexts which resist substitution, exhibit existence neutrality, or lend themselves to unspecific readings.

Here I am primarily interested in existence neutrality. I believe that perceptual verbs can exhibit all marks of intensionality, but existence neutrality is the most relevant to my project.

In fact, the kind of intensionality which interests me here is not exactly existence neutrality as glossed above. I characterized existence-neutral readings as readings on which "existential quantifiers and singular terms do not have their usual existential import". This account does not apply to a statement such as 3.3, because this statement contains no existential quantifiers or singular terms as part of the

complement of the verb.[2]

(3.3) I ordered cars.

Yet 3.3 seems susceptible to the same kind of ambiguity as "I ordered a car". More specifically, there is a reading of 3.3 on which it entails that there are cars I ordered (the material reading), and there is another reading on which it has no such implication (the intensional reading). On the intensional reading, but not on the material reading, the statement could be true even if there were no such things as cars. Likewise, a statement such as 3.4 can be read in at least two ways which parallel the material and intensional readings of "I ordered a car".

(3.4) Bob ordered every new car.

On its material reading, statement 3.4 would be vacuously true if there were no new cars in the domain of discourse, because it says merely that every new car is such that it has been ordered by Bob. By contrast, statement 3.4 would not automatically be satisfied on its intensional reading if there were no new cars, because this reading entails that Bob has performed an act of a certain kind, and the fact that there are no new cars does not entail that Bob has performed such an act.

These examples point toward a broader kind of intensionality than existence neutrality as defined above, and this kind of intensionality is more directly relevant to my aims here. This more general phenomenon is best characterized in terms of quantifier scope. Some statements of the form "α ϕ-s S" can be read either materially or intensionally. On a material reading, the argument of ϕ specified by S is or involves a variable bound by a quantifier located outside the arguments of ϕ. On an intensional reading, the proposition expressed involves no quantification of this kind.

[2]Of course, the statement does contain a quantifier according to the widely accepted view that all NPs are generalized quantifiers (Westerståhl 2008), namely "cars". But it is unclear how to extend the notion of an existential quantifier to all relevant generalized quantifiers consistently with the present account of existence neutrality: we cannot say that existential quantifiers support existential generalization, because existence-neutral contexts undermine existential generalization.

Material reading A reading of an expression of the form "α ϕ-s S", where ϕ is a verb and S complements ϕ, is material just in case its assigns the expression a logical form in which the argument of ϕ specified by S is or involves a variable bound by a quantifier located outside the arguments of ϕ.

Intensional reading A reading of an expression of the form "α ϕ-s S", where ϕ is a verb and S complements ϕ, is intensional just in case it is not material.

For the purposes of illustration, let us assume that names contribute individuals to logical form. Given this assumption (and setting aside complications pertaining to tense), the material reading of "Alice ordered a painting" ascribes it the following logical form:

(3.5) $\exists x(painting(x) \wedge ordered(Alice, x))$

This is a material reading because the second argument of the relation specified by "a painting" is a variable bound by a quantifier outside the arguments of the relation. By contrast, the intensional reading of this expression assigns it a logical form in which the second argument of the relation is not (and does not involve) a variable bound by a quantifier outside the argument. On this alternate reading, the statement does not entail that there is any painting. We will see in the next chapter what is the logical form of such a statement on its intensional reading. I want to leave all options open for now. Multiple possibilities will be assessed in the next chapter.

A potential flaw of the present definition of intensional readings is that some statements which do not intuitively seem intensional in any way might turn out intensional. In particular, a statement such as "Alice sees Bob" might be thought to have only one possible reading: $see(Alice, Bob)$, which involves no quantifiers at all. If this is correct, "Alice sees Bob" is always intensional on my definition, which might seem inconsistent with the normal usage of "intensional".

This issue could be addressed by stipulating that a proposition of the form "ϕ(___, α)", where α is an individual or an entity constituted in part by an individ-

ual, always constitutes a material reading. Having said this, I prefer the unmodified definition because I believe that perceptual ascriptions such as "Alice sees Bob" are capable of material readings even on this definition, and that the material readings are in fact the most commonly intended ones. On broadly descriptivist views of the logical role of names, they contribute the logical equivalent of (possibly rigidified) definitive descriptions to propositions. They are therefore quantificational expressions. On such a view it is natural to allow the quantifier contributed by "Bob" in "Alice sees Bob" to take either wide scope or narrow scope over the argument of "see". The wide scope reading is material, while the narrow scope reading is intensional. On a broadly descriptivist view of names, then, it is not true that all readings of "Alice sees Bob" are intensional. I'll discuss some evidence for this position in section 4.6 after having introduced necessary background material. We will see at the same time that the understanding of the intensional / material distinction I am proposing is in line with the leading account of intensional transitives in Montague semantics.

While my definition of intensional readings is not meant to capture all features commonly attributed to intensional contexts, it worth noting that intensional readings as defined above are not unrelated to non-specific readings and substitution resistance. Indeed, intensional readings seem to be exactly the same as non-specific readings on my definition of the former. They also seem to coincide with the failure of substitutivity of co-extensive terms as far as definite descriptions go. When a definite description takes wide scope over a predicate, the resulting logical form involves an existential quantifier outside the argument of the predicate. The resulting reading is therefore material. Co-extensive definite descriptions with wide scope are also arguably interchangeable *salva veritate*. On the other hand, definite descriptions which have narrow scope do not result in intensional readings and are arguably not as a general rule interchangeable *salva veritate*. So intensionality and the failure of substitutivity of co-extensive terms go together as far as predicates

involving definite descriptions and no other quantificational expressions go.

I will refer to the states ascribed on intensional readings of perceptual ascriptions as *intensional perceivings.*

Intensional perceiving A state ascribable using a locution of the form "α ϕ-s S", where ϕ is a perceptual verb used intensionally and S complements ϕ.

The perceptual ascriptions with which I began this section are all naturally read as ascribing intensional perceivings in this sense.[3]

A few claims I am *not* making need to be flagged before continuing. I claim that it is typically possible to read perceptual ascriptions in an intensional manner, but I do not claim that this is always possible. For example, I suspect that the statement "I see the man next door" has no (meaningful) intensional reading. This is fine for my purposes. I also do not claim that the intensional reading of "see" is a reading on which it can be true for any property F that one sees an F without there being an F. For example, I cannot see_i an object without there being an object, because my seeing an object requires that I (an object) exist. The fallibility of intensional seeings is limited just like the fallibility of belief (necessary propositions and contingent propositions such as that something is believed cannot be believed falsely). It is precisely this observation which led me to think of the distinction purely in terms of logical form and quantification.

Ayer (1940) draws a distinction which is superficially similar to my intensional / material distinction:

> If I say that I am seeing a stick which looks crooked, I do not imply that anything really is crooked ... or if, being subject to an illusion of double vision, I say that I am perceiving two pieces of paper, I need not be implying that there really are two pieces of paper there. But surely,

[3] Anscombe's burning rubber example might be an exception: the definite description "the smell of burning rubber" is more naturally read as having wide scope. But one could say "I smell smell 34" instead, where smell 34 is the smell of burning rubber.

> it may be said, if the two pieces of paper really are perceived they must both exist in some sense, even if not as material things. The answer to this objection is that it is based on a misunderstanding of the way in which I am using the word "perceive". I am using it here in such a way that to say of an object that it is perceived does not entail saying that it exists in any sense at all. And this is a perfectly correct and familiar usage of the word. (Ayer 1940: ch 1)

Ayer seems to want to highlight the possibility of intensional readings of perceptual ascriptions in something like my sense. This passage is followed by a summary of the main features of the two kinds of reading of perceptual ascriptions he wants to distinguish. In one sense of "see", Ayer claims, "it is necessary that what is seen should really exist, but not necessary that it should have the qualities that it appears to have." In another, "it is not possible that anything should seem to have qualities that it does not really have, but also not necessary that what is seen should really exist." Ayer's distinction appears to turn on what I call intensionality and what we might refer to as *appearance neutrality*. It is similar to the intensional / material distinction I have drawn, in that, at least as far as perceptual predicates of the form "ϕ-s an F" go, it separates readings which are existence-neutral but are not appearance-neutral (intensional readings) from others which are not existence-neutral but are appearance-neutral (material readings).

Having said this, there are important differences between Ayer's characterization of these features of perceptual ascriptions and mine. First, Ayer's characterization of the existence neutrality of intensional readings seems to carry a commitment to Meinongian objects. He appears to be saying that it is (in practice) possible for there to be an object x such that one can correctly say of x that it is perceived without its existing ("to say of an object that it is perceived does not entail saying that it exists in any sense at all"). This claim commits one to there being non-existent objects. By contrast, my characterization of intensional readings does not: I am

only committed to it being (in practice) possible to see an F without there being an F. All that this requires is that the determiner occurring as part of the complement of "see" not take wide scope over the latter, as in “I requested a diamond table”. I will explain how this works in section 4.6.

Ayer’s characterization of appearance neutrality (or non-neutrality) also differs significantly from mine. His claim that there is a sense of "see" on which "it is not possible that anything should seem to have qualities that it does not really have" is rather puzzling. This is supposed to be a feature of intensional readings, those readings which exhibit something like existence neutrality. This claim is puzzling because when one says that one "sees$_i$ an F", it is unclear that one can be construed as talking about any object one sees. It is consequently unclear how to understand Ayer’s reference to the "it" that seems a certain way in such a state. My characterization of appearance neutrality does not have this shortcoming, because it is not stated as a relation between the object one perceives and the manner in which it appears: I merely said that intensional readings but not material readings carry substantive implications regarding how the world appears to one.

These differences between Ayer’s formulation and mine make all the difference when it comes to Austin’s criticisms of Ayer. Take first this passage from *Sense and Sensibilia*:

> ... you may remember that he [Ayer] said earlier, as explicitly as could be, that there is a "correct and familiar" usage of "perceive" which is such that "to say of an object that it is perceived does not entail saying that it exists *in any sense at all*". On this there is no possible comment except that there isn’t. (Austin 1962: 95)

Austin is right: it is not possible for there to be an *x* which does not exist. But this passage does not contain any argument against intensional readings of perceptual ascriptions as defined here.

Austin criticizes other features of Ayer's account which are not part of the intensional-material distinction I am proposing. In particular, a number of Ayer's examples are flawed, and Austin correctly points this out. Austin also objects to details of Ayer's characterization of the appearance neutrality of material readings which are not part of my account. Aside from the above passage, the only part of Austin's discussion which might reasonably seem to carry over to my characterization of the intensional / material distinction is the part where he insists that Ayer has not revealed a distinction between two *senses* of "perceive" but something else of lesser importance—that we can "stretch" the "ordinary usage" of perceptual verbs in exceptional circumstances:

> Since, in this exceptional situation, though there is only one piece of paper I seem to see two, I may want to say "I am perceiving two pieces of paper" *faute de mieux,* knowing quite well that the situation isn't really that in which these words are perfectly appropriate. But the fact that an exceptional situation may thus induce me to use words primarily appropriate for a different, normal situation is nothing like enough to establish that there are, in general, two different, normal ("correct and familiar") *senses* of the words I use, or of any one of them. To produce a rather baffling abnormality like double vision could establish only, at most, that ordinary usage sometimes has to be stretched to accommodate exceptional situations. (Austin 1962: 91)

Here again, Austin is picking on a relatively superficial feature of Ayer's exposition. Ayer does imply at times that there are two senses of "perceive", but this is not part of the intensional / material distinction. What matters is that perceptual predicates such as "ϕ-s an F" can be read as ascribing two different kinds of state (intensional and material perceivings). This can be the case without perceptual verbs being ambiguous. Indeed, the definitions of intensional and material readings I gave nearly imply that the intensional / material distinction lies in an ambiguity in logical form

rather than a lexical ambiguity. In this regard, my approach is in line with leading accounts of the intensional / material ambiguity (see section 4.6).

One point which must be conceded is that intensional readings are atypical, in that material readings are usually intended. But this does not mean that they are not available as part of ordinary language. There is a simple explanation of their atypicality. First, our aim when talking about what we "perceive" is normally to communicate what there is in the environment, and only material readings encode this. There is also the fact that it is typically harder to tell what one perceives in the intensional sense than it is in the material sense, because intensional readings capture aspects of how the world appears to one "subjectively". It is comparatively easier to determine what objects one is being perceptually affected by. So material readings are usually both less committal and more to the point. This goes a long way toward explaining why they have default status in an ordinary conversational context.

3.2 Intensional perceivings and phenomenology

My aim in this section is to clarify the relation between intensional perceivings and phenomenal states.

3.2.1 Phenomenal uses of perceptual verbs

Intensional perceivings are generally accompanied by experiences. Normally, when one sees_i an F, one also experiences an F. For example, normally when you see_i a flash of light, you experience a flash of light; normally, when you see_i a red expanse, you experience a red expanse; normally, when you see_i a tree shape, you experience a tree shape; and so on. The same is true of intensional perceivings in other modalities. For example, normally, if you feel_i a sensation in your back, you experience a sensation in your back.

On the other hand, we sometimes ascribe intensional perceivings in absence of corresponding phenomenal states. For example, if I say that Bob sees thieves everywhere, I don't mean to suggest that Bob instantiates phenomenal states of any kind. Or one might say that one can hear Walter Cronkite talk even though Cronkite isn't there and the phenomenology of one's experience at the time involves nothing Cronkite-specific. This could perhaps be the case in virtue of one's taking a certain (apparent) voice to be the voice of Walter Cronkite. There might also be a sense of "feel$_i$" which implies no phenomenology. For example, it might be that "Bob feels$_i$ a threat" does not imply that Bob instantiates any phenomenal state.

While some intensional perceptual ascriptions might not be intended to ascribe phenomenal states, we can distinguish a kind of reading of intensional perceptual ascriptions on which they are equivalent to "experiencing" ascriptions in our technical sense. In the *phenomenal sense*, to say that one ϕ-s S, for any expression S which complements the perceptual verb ϕ, is simply to say that one experiences S in the relevant modality. To see a red ball in the phenomenal sense (see$_p$), for example, is simply to visually experience a red ball.

The following story illustrates the phenomenal use of "see". Imagine that you have just undergone surgery to receive a cerebral implant that is supposed to augment your visualization abilities. A nurse will perform some tests on you to verify that the implant functions properly. She begins by attaching a device on your scalp.

– Now tell me if you see a blue grid.

She flips a switch, which results in your experiencing a green but slightly bluish grid for the briefest moment. Having been primed to expect a blue grid, you take it to be blue.

– Yes, I saw a blue grid.

The nurse initially seems satisfied by your response, but her expression quickly changes into a frown as she takes a second look at her console.

– I got the settings wrong. You've really seen a blue grid?

– Yes I have.

– Something is not right. I will run a scan of your short-term visual memory.

A moment later...

– Thanks God you haven't really seen a blue grid! Why on Earth did you say you had? You don't know the difference between green and blue?

– Oh, well, sorry. I really thought it was blue. Doesn't that count as seeing it blue, anyway?

– No, it doesn't.

It is natural to take the nurse to be asking whether you saw_p a blue grid (whether you visually experienced a blue grid). This example at least makes clear that there is a sense of "see" on which seeing something involves visually experiencing it.

One might think that something else than experience is involved in this kind of seeing, so it is worth considering the options. What else could the nurse have meant when she asked if you saw a blue grid in addition to your experiencing a blue grid? The nurse did not want to know whether you thought (believed, judged, etc.) that there really was a blue grid in front of you, because she knew that there was not and did not expect to convince you that there was. She did not want to know about the external cause or normal cause of the state you were in either. She knew about its external cause already, and insisting that the state you were in is normally caused by blue grids would not satisfy her. She might reply: "for all I know you're one of those inverts—I want to know how things *appeared* to you". She was also not asking about your physiological condition. Indirectly, perhaps, she was (if physicalism is true), but her meaning was not "are you in such and such brain state", because she

knows that you would not have been able to answer such a question. So, seeing a blue grid in the nurse's sense does not conceptually require that some condition pertaining to judgments, the causes of the experience, or one's physiology obtains. It also does not seem to be a matter of having an experience of a certain kind while a disjunction or other logical combination of these conditions obtains. There seems to be nothing else for the nurse to be asking about but phenomenology. Consequently, it seems that by "see a blue grid", she meant "visually experience a blue grid". In general, when we speak of seeing$_i$ things of the sort we can perceive independently of judging that things are thus and so, we mean seeing in the phenomenal sense. Seeing$_i$ flashes, grids, colors, shapes, etc, on the default understanding, are states of visually experiencing flashes, grids, colors, shapes, etc.[4]

What goes for sight goes for other modalities. Suppose for example that you are hearing a ringing in your ears. Typically, part of what you mean when you say that you are hearing a ringing is that you are having an experience of a certain kind, namely, an aural experience of a ringing in your hear. What else might you intend to say? Normally, to say that you are hearing a ringing is not to say anything about your beliefs, your environment, your physical state, or the current or normal circumstances of your state, because you could conceivably hear a ringing whatever your beliefs or the state of your body or environment. As in the case of blue grids, there seems to be nothing for the expression "hearing a ringing" to pick out but a certain experiencing state. Likewise for other modalities. If I say that I am feeling pain in my phantom arm, for example, all I am talking about is the experience I am undergoing.

Phenomenal readings of perceptual ascriptions can also be brought out by reflection on the grounds of certain ascriptions. Suppose that you get up from your bed too quickly and "see stars" as a result of this. You say, "Oh! I see stars". That you are seeing stars is something that you would normally realize spontaneously.

[4]Jackson (1977) and Chisholm (1948; 1950) distinguish phenomenal readings of "look" statements in this way.

As a result, it is plausible that you would normally come to notice that you are seeing stars either perceptually or introspectively: normally, your ground for asserting that you see stars is either a perceptual judgment or an introspective judgment. Now, you do not normally make any relevant perceptual judgments when you find yourself seeing stars. As a general rule, you know full well that there are no specks of light in front of you. So it seems that the basis of your assertion, in the normal case at least, must be an introspective judgment. It is plausible that there are only two kinds of mental state we can introspect: attitudes and phenomenal states. Since seeing$_i$ stars does not seem to involve having an attitude or making a judgment about stars,[5] this makes it seems likely that it consists in nothing more than a phenomenal state. The same line of argument straightforwardly applies to other intensional phenomenal ascriptions.

In addition to illustrating a phenomenal use of a perceptual verb, the preceding example illustrates an important observation regarding such uses. It is clear that when one says that one sees$_p$ stars, one does not mean that one is perceiving$_p$ celestial bodies. The word "stars" here functions as a shorthand for something like "multiple brilliant specks of lights". The use of such shorthands in phenomenal perceptual ascriptions is common and can obscure the distinction between phenomenal and non-phenomenal ascriptions. If I say that I saw a car, an elephant, or an alien while hallucinating, for example, I might seem to be using "saw" in some non-phenomenal intensional sense, because I arguably cannot experience a car, an elephant, or an alien (in the technical sense of "experience"; I can in the everyday sense). But this is generally not the case. Typically, someone who claims to have seen a car, an elephant, an alien or any other "high level" entity in a hallucinatory context does not intend the ascription to be taken literally. If I report that I saw a pink elephant in my lounge while under the effect of LSD, I don't mean that I saw a member of the taxonomic family of elephantidae in any sense of "saw". I mean

[5] See Bourget (forthcoming).

that I saw_{p} something pink with certain readily perceptible features of elephants. Similarly, when I say that I saw a car in my dream, I don't mean to imply that I saw an artifact whose primary function was to transport people and goods in such and such manner. I mean that I saw_{p} something with certain superficial car features. We often use loaded terms like these as shorthands for more accurate descriptions of what we see in the phenomenal sense, and this practice can create an illusion that we mean "see" in some other sense. I have used and will continue to use such shorthands for legibility's sake, but it is important to keep this in mind.

Of course, it might be that we can literally see_{p} many different kinds of high-level entity such as artifacts and life forms. We don't have to take a stand on this matter for present purposes. The point is that, even if we suppose that we cannot experience such entities, we should not take the fact that most candidate phenomenal uses of perceptual verbs in real life involve high-level terms to suggest that there are few or no such uses.

While intensional and phenomenal readings of perceptual ascriptions generally go together, the intensional / material distinction is in fact orthogonal to the phenomenal / non-phenomenal distinction. Contrary to the intensional / material distinction, the phenomenal / non-phenomenal distinction appears to be a distinction between two kinds of understanding of "see", as opposed to the logical role of surrounding expressions in the statement. For example, Bob does not see thieves in the same sense in which one can see color patches. Seeing in the first sense is primarily a cognitive or epistemic phenomenon (perhaps one that involves some phenomenology, but it is not the same as experiencing in the technical sense). Let us mark non-phenomenal readings with a "c" subscript (for "cognitive" or "causal", labels which arguably characterize the two main kinds of non-phenomenal reading). All of the following combinations of subscripts are valid:

(3.6) I see_{ip} a round shape

Entails that I am experiencing a round shape, but not that there is a round

shape.

(3.7) I see_{mp} a popular colored shape

Entails that there is a colored shape which is popular, and that I am experiencing it.

(3.8) I see_{ic} a bad omen

Entails that I take there to be a bad omen, but not that there is one.

(3.9) I see_{mc} a bad omen

Entails that there is a bad omen, and that I am aware of it (but perhaps not as such).

Material-phenomenal readings of perceptual ascriptions ascribe more than phenomenal states. Readings of the type required by statement 3.7, in particular, do not ascribe states which are individuated solely by what it is like to be in them. Statement 3.7 ascribes a property of roughly this form:

$$\lambda x(\exists y(popularColoredShape(y) \wedge see_p(x,y)))$$

This property is not a phenomenal state even though its instantiation requires having an experience of a certain general kind (an experience of a colored shape). By contrast, statement 3.6 plausibly ascribes a pure phenomenal state.

The availability of "mp" readings of perceptual ascriptions can be traced back to the availability of material readings of "experiencing" ascriptions. A statement such as 3.10, for example, naturally lends itself to a material reading along the same lines as statement 3.7.

(3.10) I am experiencing a popular colored shape.

This statement does not ascribe a pure phenomenal state on its material reading. So while there are uses of perceptual ascriptions which are equivalent to "experiencing" ascriptions, the latter do not all ascribe pure phenomenal states, because they are

themselves subject to the material / intensional ambiguity. To ascribe a phenomenal state and nothing more using a perceptual verb, one must use it both phenomenally and intensionally.

I am going to refer to perceptual ascriptions which ascribe phenomenal states (and nothing else) as *pure phenomenal ascriptions.*

Pure phenomenal ascription A perceptual ascription of the form "NP ϕ S" in which the predicate "ϕ S" denotes the state denoted by "ϕ-ly experiencing$_i$ S".

I will refer to the states ascribed by pure phenomenal ascriptions as *phenomenal perceivings.*

Phenomenal perceiving A state ascribed by a pure phenomenal ascription.

We can think of phenomenal perceivings as phenomenal states which can be referred to as states of ϕ-ing certain things (where ϕ is a perceptual verb).

I will for the most part leave out "i" and "p" subscripts from now on. Intensional and phenomenal readings should be assumed unless the context clearly indicates otherwise.

3.2.2 The perceptual conception of sensory experience

We can now state the perceptual conception of sensory experience more precisely.

The perceptual conception of sensory experience (PCSE) For every sensory phenomenal state *e* in modality ϕ, there is a phenomenal perceiving *i* in modality ϕ such that $e = i$.

My case for PCSE starts from the observation that all sensory phenomenal states can *in principle* be designated using constructions of the form

ϕ-ly experiencing$_i$ S,

where ϕ specifies a sensory modality and S complements the verb. My argument for PCSE goes as follows:

The argument for PCSE

1. Every sensory phenomenal state in modality ϕ is a state of ϕ-ly experiencing$_i$ something in modality ϕ.

2. Every state of ϕ-ly experiencing$_i$ something in modality ϕ is a phenomenal perceiving in modality ϕ.

   Therefore, for every sensory phenomenal state e in modality ϕ, there is a phenomenal perceiving i in modality ϕ such that $e = i$.

Philosophers writing on consciousness systematically designate phenomenal states as states of ϕ-ly experiencing certain things. For example, we say that visually experience colors and shapes, or that we aurally experience sounds. We can also describe more complex states using the same language, e.g. the state in which one visually experiences a bright surface framed by a dark region and punctuated with little black marks. Premise 1 therefore seems prima facie plausible. We do not normally describe bodily experiences as states in which we bodily experience certain things—normally, we only describe them as states in which we experience certain sensations. Still, the "bodily" qualification does not hurt, and it seems possible to describe bodily experiences as states in which we bodily experience certain things.

While premise 1 is prima facie plausible, there are two major potential exceptions to it.

First, disjunctivism might reasonably be taken to imply that premise 1 is not true of many (if not all) phenomenal states. Recall that disjunctivists deny that the same phenomenal states can occur in veridical and non-veridical perception. If this were correct, it would seem that one must explicitly or implicitly specify whether a state is of the veridical or non-veridical kind in order to successfully pick it out, and expressions of the form "ϕ-ly experiencing S" do not seem to do that.

Let us suppose that states of veridically and non-veridically experiencing an F are distinct phenomenal states as the objection requires. Then the question arises

as to which kind of experiencing phenomenal perceivings are, because we know already that phenomenal perceivings are phenomenal states and experiencing states as well. There are three possibilities: all phenomenal perceivings are non-veridical experiencings, all phenomenal perceivings are veridical experiencings, or some phenomenal perceivings are veridical experiencings while some are non-veridical experiencings.

The first possibility might initially seem plausible since our main examples of intensional and phenomenal perceivings are all illusory or hallucinatory. But the only reason I had to focus on non-veridical cases is that the contrast between perceiving$_{i}$ and perceiving$_{m}$ is less obvious in veridical cases, because, normally, one simultaneously perceives$_{i}$ and perceives$_{m}$ in veridical cases. Prima facie—absent independent evidence for disjunctivism of the kind I'll discuss later—, there does not seem to be any reason to think that phenomenal perceivings occur only when what one perceives in them does not obtain. The state of seeing$_{ip}$ stars, for example, could very well be instantiated by someone who really is facing brilliant specks of light. So it seems that seeing$_{ip}$ stars is something that one can do whether there are stars or not, in the relevant sense of "stars". Or consider this case reported by Treisman and Schmidt (1982). Apparently, subjects presented in rapid succession with a pink X and a yellow T reliably "see" a solid, steady pink T. Illusions of this kind plausibly involve phenomenal perceivings which are phenomenally identical to actual veridical phenomenal states. Disjunctivists deny that this is possible, but one can only deny this possibility based on independent evidence for disjunctivism; prima facie, the possibility is genuine. This example and the example of seeing$_{ip}$ stars seem to me to constitute a good prima facie case against all three of the aforementioned possibilities.

It also helps to reflect on how perceptual verbs are in fact used. Imagine for example that you are hallucinating a pink elephant in your lounge. As a result, you say, "I see a pink elephant", meaning of course "see$_{ip}$". If you are like the aver-

age person, learning that there is a pink elephant in front of you (or that there is not) would never cause you to take back your statement. I am here merely stating a fact about how "seeing" locutions are used, not proposing a far-fetched thought experiment. It is a fact that laymen describe their phenomenal states using perceptual verbs independently of what there is to perceive around them, and this strongly suggests that phenomenal perceivings are independent of one's environment.

Another potential source of counterexamples to premise 1 is experiences which reflect certain kinds of perceptual variation, e.g. blur, perspective, and double vision. A phenomenal state which reflects perspective in its phenomenology cannot readily be designated as a state of visually experiencing X or a state of having a visual experience of X, for any X. Similarly, it is not obvious at first how to capture the phenomenology of blurry or double visual experiences by filling in the X.

In response to this, notice that the phenomenology of these experiences is hard to capture *simpliciter*. This is most evident in the case of perspective. I don't know how to intrinsically characterize the phenomenal character of a visual experience in a way that would capture the effect of perspective. I can say that I am having an experience of an F of the kind that one has when one looks at an F from such and such viewpoint, but I cannot directly characterize the phenomenal character of the experience, whether using an "experiencing ___" locution or not.

This suggests that the reason why experiences which reflect the effects of perceptual variation—the tricky ones often brought up as putative counterexamples to representationalism, in any case—are not readily describable as visual experiencings of certain entities is that we have an insufficient grasp on what is to be described in this way, and not that this cannot be done in principle. One should at least agree that the nature of these experiences is elusive conceptually. As a result, proper methodology recommends that we extend to them whatever theory works best for other states. We should take simpler phenomenal states such as *experiencing a red square*, see what is true of them, then generalize to the more complex cases. Since

most simple experiences can be described as states of ϕ-ly experiencing certain things, it seems a reasonable strategy to tentatively suppose that all can in principle. I will return to perceptual variation in chapter 6, where I will supply more evidence to show that the relevant experiences do not make exception to PCSE or virtualism. For now, I move on to the second premise of the argument for PCSE.

Let F be the function which maps states of ϕ-ly experiencing certain things to phenomenal perceivings as follows, for any phrase S which can complement the verb to experience:

visually experiencing$_{\text{i}}$ S $\Rightarrow$ seeing$_{\text{ip}}$ S

aurally experiencing$_{\text{i}}$ S $\Rightarrow$ hearing$_{\text{ip}}$ S

bodily experiencing$_{\text{i}}$ S $\Rightarrow$ feeling$_{\text{ip}}$ S

...

For example, the value of F for the state of visually experiencing$_{\text{i}}$ a green square is the state of seeing$_{\text{ip}}$ a green square. The value of F for the state of bodily experiencing$_{\text{i}}$ a sensation in your left leg is the state of feeling$_{\text{ip}}$ a sensation in your left leg.

Given this specification of F, we can show not only that premise 2 of the argument for PCSE is true but that the following stronger thesis holds:

Identity thesis For every state e of ϕ-ly experiencing$_{\text{i}}$ something, e = F(e).

The identity thesis is true by construction. First, F is defined for all experiencing$_{\text{i}}$ states. Second, for all e, it is true by the definitions of F and phenomenal perceivings that F(e) = e, because to ϕ S on a pure phenomenal reading is by definition to experience$_{\text{i}}$ S in the relevant modality.

There is another line of argument for PCSE which is worth mentioning briefly. In introduction we saw that the word "experience", in everyday language, can only be used to pick out phenomenal states when combined with the somewhat metaphorical "what it's like". While I hope that I have succeeded in directing your attention

to the right phenomenon by talking about "what it's like", I am under no illusion of having offered a definition of phenomenal states in everyday language. I merely gestured toward the concept. The realization that I had to engage in such hand-waving to explain what I mean by "phenomenal state" gave me serious doubts about this notion until I found that these states could be described precisely in plain language using perceptual verbs. One reason to endorse the perceptual conception of sensory experience is that if it is wrong there is arguably no way to spell out—as opposed to merely gesturing toward—the general notion of a phenomenal state in everyday language. By contrast, given the perceptual conception of sensory experience, we can say simply that a phenomenal state is a state in which one $\text{perceives}_{\text{ip}}$ something. The distinction between phenomenal and non-phenomenal perceivings still has to be explained through examples, but at least it is there to be used in everyday language. Since it is already part of the folk conceptual repertoire, it is not significantly harder to convey than the distinction between the different meanings of the word "bank".

3.3 The basic-derivative distinction revisited

In chapter 1, I drew a distinction between basic and derivative phenomenal states: the latter, but not the former, are mere states of instantiating phenomenal states distinct from themselves. I will conclude this chapter by relating this distinction with yet another distinction between two kinds of perceptual ascription. This is the last preparatory step before we can turn to the case for virtualism in the next chapters.

All perceptual predicates which lend themselves to pure phenomenal readings have the form "ϕ-ing S" or "ϕ-s S", where S complements the verb. In some pure phenomenal ascriptions, S is a bare infinitive clause, a participial clause, or a small clause. Here are some examples:

(3.11) She sees [a pink elephant fly].

(3.12) She hears [noise move around the room].

(3.13) She feels [the sensation intensify in her stomach].

(3.14) She sees [a round shape (be) in front of herself].

(3.15) She feels [a sensation (be) in her chest].

(3.16) She sees [a man dancing].

(3.17) She hears [a noise fading].

Statements 3.11, 3.12, and 3.13 illustrate the use of bare infinitive clauses as complements of perceptual verbs. 3.14 and 3.15 illustrate the use of small clauses as complements. Statements 3.16 and 3.17 have participial clauses as complements. These three kinds of ascription are not unrelated.[6] They are also importantly different from other kinds of ascription. I will refer to the states they ascribe as *B-perceivings.* I will refer to states which *cannot* be ascribed using bare infinitive, participial or small clause complements as *D-perceivings.*

Here are examples of perceptual ascriptions which plausibly ascribe D-perceivings while lending themselves to pure phenomenal readings:

(3.18) He sees [a pink elephant].

(3.19) He hears [some noise].

(3.20) He feels [a sensation].

(3.21) He feels [something].

All four examples involve "determiner phrases" (a sometimes controversial label) as complements of perceptual verbs. While there are many more grammatical types of

[6]Small clauses are arguably contractions of bare infinitive clauses, and participial clauses are nearly synonymous with their bare infinitive counterparts on intensional readings (though not on other readings they allow; see Declerck 1982, Felser 1999).

perceptual ascription, arguably only ascriptions of the "ϕ DP" type and those which define B-perceivings lend themselves to pure phenomenal readings. I will say more about potential exceptions below.

The distinction between B-perceivings and D-perceivings is important because it corresponds to a large extent to the distinction between basic and derivative phenomenal states: all basic phenomenal states are (phenomenal) B-perceivings (hence the "B" and "D" labels). I will refer to this thesis as "PCSE+":

PCSE+ For every basic sensory phenomenal state e in modality ϕ, there is a phenomenal B-perceiving i in modality ϕ such that $e = i$.

PCSE+ is the claim which will be the starting point of my argument for virtualism in the next chapter. Given PCSE, we can establish PCSE+ by showing that every basic phenomenal state which is a phenomenal perceiving is a B-perceiving. Put differently, we need to show that every D-perceiving is derivative if phenomenal.

That D-perceivings are at best derivative phenomenal states is obvious in some cases. Take for example the state of *feeling something* (3.21). This state can be realized by any number of more specific phenomenal states, and it cannot obtain without another phenomenal state realizing it. The D-perceiving *seeing a pink elephant* is also a derivative state, because this is not something that one can do by itself. To see a pink elephant, one has to see a pink elephant do something or other: fly, walk, bath, sit there, etc. The state of seeing something red also doesn't seem to be capable of instantiation on its own. In order to see something red, one need at least see something *be* red. The state of seeing something red is derivative on the wide variety of more specific B-perceivings which involve redness.

Of course, there is a question as to whether seeing$_{\text{p}}$ something be red is not also a derivative state. Perhaps one can only directly see$_{\text{p}}$ determinate shades of red, and perhaps one needs to see$_{\text{p}}$ an object as having a particular shape in order to see$_{\text{p}}$ it as having a certain color. But what matters for now is that phenomenal D-perceivings are derivative. That some phenomenal B-perceivings are also derivative is allowed

by PCSE+.

Aside from D-perceivings ascribed by predicates of the form "ϕ-s DP", there are two other types of perceptual ascription which might be thought to allow pure phenomenal readings:

(3.22) I hear that Bob has graduated.

(3.23) I hear Bob.

Perceptual predicates with "that"-clauses as complements do not lend themselves to pure phenomenal readings. Note first that only the verbs to hear, to feel, and to see accept "that"-clauses (one cannot smell that ... or taste that ...). In the case of "hear that" and "feel that", the states ascribed have plainly nothing to do with phenomenology. This is less obvious in the case of "see that", because one tends to use "see that" when one has relevant experiences (sometimes merely cognitive experiences). Still, it is perfectly acceptable to say that one sees that such and such is the case without having any relevant experiences. One might want to distinguish this use of "see that" from a more perceptual one, but I don't see where to draw the line. "Seeing that" seems to me to be a dead metaphor. The fact that "hearing that" and "feeling that" have clearly nothing to do with perception strongly suggests this.

In any case, there is another reason why perceptual ascriptions complemented by "that"-clauses do not allow pure phenomenal readings. The reason is that they are factive: seeing that P entails that P is the case. I argued earlier that phenomenal perceivings are independent of external facts. Even if it were possible to see_p that P, the state ascribed would not be independent of the fact that P. At best, seeing_p that P is a matter of instantiating a certain phenomenal state while a certain fact obtains. It is not a state individuated solely by what it's like to be in it.

We can remain neutral on whether perceptual predicates with simple noun complements (e.g. "see Bob") have pure phenomenal readings. In this case it is at least clear that *if* constructions of this kind have pure phenomenal readings, the states

they ascribe are derivative phenomenal states: if one can see_p an individual (e.g. Bob), that is at least in part in virtue of seeing_p some of its features. There is no mere seeing a bare individual in the phenomenal sense of "see". Just ask yourself what it's like to see Bob. Nothing in particular. Of course, this does not mean that it is impossible to see bare individuals in some non-phenomenal sense. Note also that this does not mean that we see_p individuals by seeing_p generic facts (e.g. by seeing something be Bob-like): it could be that seeing_p Bob (if possible at all) is grounded in such states as seeing_p Bob walk, stand still, wear a red t-shirt, ecetera, which could be phenomenal B-perceivings. (But it could also be that there is no such thing as seeing_p Bob.)

The role of time in experience provides further evidence for PCSE+. While B-perceivings appear to ascribe experiences of time, change, events, and actions, this is not true of phenomenal D-perceivings. So another way to approach the present question is to ask whether the basic phenomenal states are states in which we experience time, events, actions, etc. We tend to think of visual experiences as static snapshots which give the illusion of movement in the same way that cinematography gives the illusion of movement through rapidly changing frames, but this cannot be right: if there is an illusion of movement at the cinema, that is because there is an experience of change; the experience of change itself cannot be an illusion. If we did not experience change itself, we could at best have the illusion of experiencing change in the sense that we could be unaware of the fact that we do not experience change. But I do have a positive phenomenal impression of experiencing things in time (and positively seem to positively seem to ...). It therefore seems plausible that we experience time, change, events, and actions. Since it is hard to see how such experiences could be realized by experiences which are not experiences of time, change, events, or actions, it seems plausible that they are basic phenomenal states. This supports the view that the basic phenomenal states are B-perceivings.

The view that we experience change has had its critics, but the main basis of

the criticism has traditionally been that we cannot experience change because this would require experiences to be temporally extended (see Dainton 2008; 2000). No such constraint applies if virtualism is true. Just like a judgment can bear on a different period of time than it itself lasts, an instantaneous experience could well have as content a state of affairs which spans a certain amount of time. We should therefore not deny the phenomenological datum on this ground. Besides, it is not true that experiences would have to be extended in time in order to capture change if virtualism were false. For example, it is compatible with the sense-datum theory that experiences and sense data are subject to continuous change in the same way that an object moving continuously in space has an instantaneous velocity given by $\frac{dx}{dt}$ for one spatial dimension x.

While the view that we experience change supports the view that the basic phenomenal states are B-perceivings, it is worth stressing that the latter does not require the former. We could just as well take the most basic experiences to be B-perceivings of things merely being some way or other in a timeless manner. I mentioned the matter only as additional evidence for the view that all basic phenomenal states are B-perceivings.

It must be emphasized that PCSE+ is compatible with a number of different stances on DP-based perceptual ascriptions. In particular, PCSE+ allows that some or even all DP-based perceptual ascriptions ascribe B-perceivings. For example, it is consistent with PCSE+ that seeing [a flying pink elephant] is the very same state as seeing [a pink elephant flying]. Even if there were no phenomenal D-perceivings at all, nothing wrong would follow for PCSE+. Some of the examples used above would be incorrect if this were the case, but PCSE+ follows from PCSE on the assumption that there are no D-perceivings, so we would not need these examples. The case for virtualism would merely be simplified.

3.4 Summary

In this chapter I have argued that basic phenomenal states are perceptual states of a certain kind. I have first distinguished between intensional and material readings of perceptual ascriptions. The distinction between the two is one of logical form. I have then distinguished between the phenomenal and non-phenomenal senses of "see" and other perceptual verbs. I have argued that sensory phenomenal states can all be picked out using perceptual verbs intensionally and phenomenally. I have further argued that the basic phenomenal states are more specifically phenomenal B-perceivings. This last claim, PCSE+, will be the starting point of my argument for virtualism in the next chapter.

Chapter 4

Sensory virtualism

In this chapter, I will argue from the perceptual conception of sensory experience defended in the preceding chapter to the sensory virtualist view (SV) defined in chapter 1.

Sensory virtualism (SV) For any sensory modality ϕ, there is a relation R such that: 1) For any basic phenomenal state s in ϕ, there is some proposition P such that s = standing in R to P; 2) R is virtual with respect to basic phenomenal states.

I will refer to claim 1 as *sensory propositionalism.* This chapter focuses almost exclusively on sensory propositionalism, because sensory propositionalism warrants claim 2 in light of the considerations in section 3.2.2 of the preceding chapter.

SV is a theory of sensory phenomenal states, not a general theory of consciousness like virtualism, so I will only talk about sensory phenomenal states in this chapter. I will continue to use the terms "phenomenal state" and "experience" without qualification to refer to sensory phenomenal states and experiences.

4.1 From PCSE+ to sensory propositionalism

The point of introducing the perceptual conception of sensory experience in the preceding chapter was that we have a much better pre-theoretic handle on what seeings and other perceivings are than "experiences". We at least have a much better handle on the semantics of perceptual verbs than we have in the case of the technical terms "experience" and "phenomenal state". Since the only way that I know of persuading someone of the correctness of sensory virtualism if they do not find it introspectively self-evident is to appeal to considerations which have to do with the semantics of the terms used to describe phenomenal states, I wanted to secure the best language possible—the language of phenomenal perceivings.[1]

My overall argument for sensory propositionalism goes as follows.

The overall argument for sensory propositionalism

1. For every basic sensory phenomenal state e in modality ϕ, there is a phenomenal B-perceiving i in modality ϕ such that $e = i$. (PCSE+)

2. For every modality ϕ, there is a relation R such that, for every phenomenal B-perceiving i in ϕ, there is a proposition P such that i = standing in R to P.

3. If (1) and (2), sensory propositionalism is true.

   Therefore, sensory propositionalism is true.

Premise 1 has been defended in the preceding chapter. Since premise 3 is trivial, I will concentrate on premise 2. I take it that premise 2 is true if the following is true:

> The proposition expressed by an ascription of a phenomenal B-perceiving of the form "α ϕ-s S" has the form $\phi(\alpha, P)$, where P is a proposition.[2]

[1] Having said this, the argument to follow could probably be recast in the "experiencing" language, at the cost of slightly weaker linguistic intuitions. If you are not convinced about PCSE+, substitute all perceptual predicates for "experiencing" predicates in what follows.

[2] One might balk at the assumption that the nature of consciousness can be inferred from the

Before we delve into the case for premise 2, it is worth noting that it is perfectly compatible with sensory propositionalism that phenomenal perceivings cannot be glossed using "that"-clauses (as noted in section 3.3). There are ways of stating relations to propositions or states of affairs which don't use explicit "that"-clauses. For example:

(4.1) I want [to learn Spanish].

(4.2) I tried [taking the bus].

The clausal complements in these sentences are naturally read as designating propositions or states of affairs: I want $learn(I, Spanish)$, I tried $take(I, bus)$. Yet "that"-clause glosses do not sound right:

(4.3) * I want that I learn Spanish.

(4.4) * I tried that I be taking the bus.

Even though "that"-clause glosses do not work with all of them, constructions of these types are commonly analyzed as ascribing relations to propositions (see for example Bealer & Monnich 1981, p. 164).

Another objection need to be addressed at the outset. One might say that phenomenal B-perceivings cannot be relations to propositions because propositions are abstract objects and abstract objects cannot be seen (because they cannot affect us causally). One response to this objection is that our understanding of propositions

logical form of phenomenal ascriptions. I believe that this assumption is harmless for two main reasons. First, it is unclear what the relevant facts about the internal structure of phenomenal states might be if they are not facts about logical form. Second, it does not really matter if all we can show is that phenomenal ascriptions have a certain logical form. The principal aim of the virtualist framework is to provide something close to a logic of phenomenal ascriptions: a precise language in which to describe phenomenal states and their inter-relations, and some ground rules capturing important generalities about phenomenal states. Even if the virtualist thesis had to be limited to a claim about the logic of phenomenal ascriptions, none of its applications would be compromised.

Note also that I am here helping myself to an account of the semantics of NPs which I will later reject. In light of what I say in section 4.6, we should allow that "α" contributes a generalized quantifier to the proposition expressed, but this does not affect the point that the ascription expresses a relation between something and P.

is sufficiently broad to include entities capable of bearing causal connections to other things. For example, we count scenes, instantiations of properties, and states of affairs as propositions. An alternative response is that it is false that all seeings require causal connections to relevant objects. In particular, it seems plausible that phenomenal and/or intensional seeings do not require such connections.

I will now state a simple argument for premise 2 of the argument for sensory propositionalism.

4.2 A quick argument

Suppose that I am undergoing a strange hallucination. I could very well say:

(4.5) What I see_p is true / false.

(4.6) What I see_p is likely to happen / unlikely to happen.

(4.7) What I see_p is possible / impossible.

4.5-4.7 seem to trivially entail 4.8-4.10.

(4.8) There is an *x* such that I see_p *x* and *x* is true / false.

(4.9) There is an *x* such that I see_p *x* and *x* is likely to happen / unlikely to happen.

(4.10) There is an *x* such that I see_p *x* and *x* is possible / impossible.

4.5-4.7 do not directly ascribe phenomenal B-perceivings, but if phenomenal B-perceivings are the basic phenomenal states as PCSE+ implies, 4.5-4.7 can only be made true by my instantiating certain phenomenal B-perceivings. Now, given the overtly dyadic use of "see" in 4.8-4.10, it seems clear that 4.5-4.7 can only be made true by my instantiating phenomenal B-perceivings which have a relational structure, i.e. states of standing in the seeing_p relation to certain entities. Moreover, it seems that the entities I stand in the seeing_p relation to in these states must be

propositions. For what else could have the properties ascribed to x in 4.8-4.10? We also say that events are unlikely to happen or are impossible, but events make perfectly fine propositions on our account of propositions.[3] We also say that sentences are true or false, but we can stipulate that this is a case where I am not talking about hallucinating sentences. It seems that for statement 4.5 to be true, for example, it has to be the case that there is a proposition x such that I stand in the seeing$_{\text{p}}$ relation to x and x is true / false.

Of course, this only establishes that some phenomenal B-perceivings satisfy premise 2 of the case for sensory propositionalism (that some phenomenal B-perceivings are relations to propositions). But notice that nothing specific has been said about what I am hallucinating. It seems plausible that my statements could be true whatever I am hallucinating. If this is correct, we can conclude that all phenomenal B-perceivings are relations to propositions, because each such perceiving should serve as truth maker for at least one possible true utterance of the form of 4.5-4.7.

Inferential connections between 4.5-4.7 and arbitrary ascriptions of specific phenomenal B-perceivings also militate for a generalization of our conclusion. For example, the following argument appears to be valid:

1. I see$_{\text{ip}}$ an elephant fly.

2. It is impossible that an elephant flies.

   Therefore, what I see$_{\text{p}}$ is impossible.

Given that, as we just saw, the conclusion entails

$$\exists x(see_p(I,x) \wedge impossible(x)),$$

explaining the validity of this inference requires that we assign premise 1 the form

$$see_p(I,P)$$

[3]That might seem like an ad hoc extension of the meaning of "proposition", but instantiations of properties (a natural understanding of events) are naturally regarded as states of affairs (Wetzel 2008), and SOAs are natural candidates to play the canonical role of propositions.

and premise 2 the form

$$impossible(P).$$

Since any phenomenal B-perceiving can play the role of premise 1 in such an argument (given a suitable premise 2), it seems that all phenomenal B-perceivings are relations to propositions.

I will now discuss a lengthier line of argument for the claim that phenomenal B-perceivings are relations to propositions. First I will argue that all phenomenal perceivings have a relational structure as opposed to a monadic structure. Then I will argue that the objects of phenomenal B-perceivings (what one is related to in them) are propositions as opposed to entities of some other kind.

4.3 The relational analysis of experience

We want to know whether the proposition expressed by a statement of the form "α sees_{ip} S" has the form R(α,__) or F(α): whether it ascribes a relational property or a monadic property. Jackson(1975; 1977) and Lycan (1987) raise some serious difficulties with monadic analyzes of perceptual ascriptions. I will not repeat their arguments here. I want to consider a different problem which arises with statements which appear to quantify over what one sees, hears, smells, etc. I think this problem cuts a little bit deeper than those raised by Jackson and Lycan. I will start with seeings.

Imagine that an "experience effector" is inserted into your brain and connected to a drawing pad. An artist is drawing on the pad, and you experience what she draws. The effector makes you experience what the artist actually means to represent, not the drawing itself. The artist could rightly say:

(4.11) You see everything I draw.

(4.12) You saw something I drew.

These sentences are particularly hard to analyze for the monadist. Other difficult cases include:

(4.13) Everything I see is not the case.

(4.14) I see what you see (in my hallucination).

(4.15) There is something I don't see. (Describing an ongoing hallucination)

The problem for the monadist is that these claims, like those I discussed in the preceding section, are naturally heard as quantifying over things seen in the phenomenal sense. Setting aside niceties of tense and other complications, the propositions expressed seem to have the following forms as far as the quantifiers and properties involved in them go:

$\forall x(draw(I,x) \rightarrow see(you,x))$

$\exists x(drew(I,x) \wedge saw(you,x))$

$\forall x(see(I,x) \rightarrow \neg iscase(x))$

$\forall x(see(you,x) \rightarrow see(I,x))$ or $\exists x(see(you,x) \wedge \forall y(see(you,y) \rightarrow x = y) \wedge see(I,x))$ (alternate reading)

$\exists x \neg see(I,x)$

One can also use an unquestionably referential term as grammatical object of "see":

(4.16) I didn't see *that.* (Describing a dream by pointing at something on TV.)

Here it seems that the proposition expressed must be

$\neg saw(I,\alpha)$

where α stands for what is being pointed at.

Arguably, none of the preceding statements are pure phenomenal ascriptions (they do not merely ascribe phenomenal states), because they only allow material readings. Still, they are phenomenal ascriptions, in that they involve the verb to

see_p, not to see_c. If the only analysis of the verb to see_p which works with them is a relational one, presumably this is true of pure phenomenal ascriptions as well.

Let us think about how a monadist could analyze the preceding statements. Some might perhaps be analyzable by quantifying over monadic seeing properties. For example, 4.14 on our first reading could perhaps lend itself to an analysis such as this:

$\forall\phi(seeing(\phi) \wedge \phi(you) \rightarrow \phi(I))$

(for any property ϕ, if ϕ is a seeing property and you have it, then I have it)

But this higher-order strategy fails miserably for 4.11, 4.12 and 4.13, because in these cases it makes no doubt that something else than a seeing property is being quantified over: the artist is not drawing a seeing property (or so I stipulate).

Alternatively, one could be tempted to give the logical form of 4.11 along the lines of

* $\forall x(draw(I,x) \rightarrow see_x(you))$

or

* $\forall x(draw(I,x) \rightarrow see_{x-ly}(you))$.

It does not really matter how the x is attached to the predicate on the RHS. The problem with this is nothing subtle: the formula is only well-formed if see_x and see_{x-ly} are interpreted as predicate symbols which have nothing to do with the variable x. That is, one might as well write

$\forall x(draw(I,x) \rightarrow F(you))$.

One might think that all the monadist needs is a little bit of lambda abstraction:

* $\forall x(draw(I,x) \rightarrow (\lambda y(\lambda z(see_y(z)))(x))(you))$

But the lambda operator can only bind a free variable. If $see_y(z)$ is intended to stand for something else than $see(z,y)$, it does not contain a variable bound by λy, and the predicate has (for all we know) nothing to do with what is being drawn.

To avoid using "see" as a binary predicate, the monadist must analyze statements 4.11-4.15 along these lines:

$seeEverythingIDraw(you)$

$seeSomethingIDraw(you)$

$whatSeeIsNotTheCase(I)$

...

These are well-formed formula, and one can stipulate that their truth conditions are as they should be. However, it seems far-fetched to suppose that this really is what the statements express. This seems far-fetched because "everything", "something", and "what", have well-defined semantics: they introduce variables. Given what we know of the semantics of these expressions, it is not psychologically realistic to say that the propositions expressed by statements 4.11-4.15 have the above forms. This is something we can agree on without agreeing on the details of the logical forms of the relevant statements. Parallel remarks apply to the demonstrative in 4.16, the pronoun "I" in 4.11 and 4.12, and the pronoun "you" in 4.14: the pronouns, at least, unquestionably refer to individuals or generalized quantifiers,[4] so the above cannot give the form of the propositions expressed by these statements. This is where my examples are more problematic than those given by Jackson and Lycan.

But the worse problem for the monadist approach only comes out when we turn to the matter of inferences. Jackson (1977) and Barwise (1981) consider inferences of the following kind:

I see an F G.

Therefore, I see a G.

[4]See section 4.6 regarding generalized quantifiers.

Inferences of this type are not strictly valid. For example, if I see a fake rabbit, it does not follow that I see a rabbit. By contrast, the following seem valid:

(4.17) You see_{p} everything I draw.

I draw_{i} a man dancing.

Therefore, you see_{ip} a man dancing.

(4.18) I see_{p} everything you see_{p}.

You see_{ip} a pink elephant fly.

Therefore, I see_{ip} a pink elephant fly.

(4.19) I don't see_{p} anything you see_{p}.

You see_{ip} a rabbit.

Therefore, I don't see_{ip} a rabbit.

Explaining the validity of these inferences would seem to require that we analyze the premises relationally. For example, argument 4.17 would become:

$\forall x(draw(I,x) \rightarrow see_p(you,x))$

$draw(I,AManDancing)$

$\therefore see_p(you,AManDancing)$

What *AManDancing* stands for will be discussed in the next section. Whatever it is, it must be one of the things that the variable in the major premise ranges over, and "I see a man dancing" has to be analyzed as ascribing a relation to this thing. Otherwise it is not possible to explain the validity of the inference.

These inferences are not only a problem for the monadist on their own. They also license the claim I made earlier that see_{p} should be analyzed relationality both on material and on intensional readings. Since any statement of the form "I see_{ip} S" can validly be inferred through arguments of the above form, all such statements must be analyzed relationally.

The argument I have presented in this section so far extends to other modalities. Take for example the following statements.

(4.20) I heard the same thing as you (but we were both hallucinating).

(4.21) I hear everything you describe (but I know there is no sound).

(4.22) I feel what you feel in your arm (but I only have a phantom arm).

(4.23) I can feel what you describe (but I know there is nothing there).

All seem to involve quantification over things perceived, and all can serve as premises of valid arguments which parallel the above. For example:

(4.24) I heard the same thing as you.

You heard noise coming from the left.

Therefore, I heard noise coming from the left.

(4.25) I can feel what you describe.

You describe a painful sensation.

Therefore, I can feel a painful sensation.

In this section I have argued that phenomenal perceivings are states of standing in certain relations to certain entities. When one says that α ϕ-s S (where the perceptual verb ϕ is used phenomenally and intensionally), one is saying that α stands in the ϕ-ing relation to a certain entity.

4.4 The objects of experience

We have seen that phenomenal perceivings are relational states. It remains to see if phenomenal B-perceivings are more specifically relations to propositions. While

the relational analysis applies both to B- and D-perceivings, I don't think that D-perceivings are relations to propositions, so I will only talk about B-perceivings in this section.

There are many candidate objects for B-perceivings, but we can divide them up into a few large classes. The relata of B-perceivings could be: a) individuals (including ordinary objects, tropes, and events qua particulars), b) properties of individuals (which I will just call "properties"), c) sets of properties of individuals, d) propositions (including SOAs and certain kinds of event), or e) entities of other kinds. I will discuss each possibility in turn.

Take first the suggestion that what we see_p in phenomenal B-perceivings are individuals. This can be cashed out in one of two ways:

I see_ip an F $\psi = \exists x(F(x) \wedge \psi(x) \wedge see_p(I,x))$

I see_ip an F $\psi = see_p(I,\alpha)$, where α is an individual to be specified

("an F ψ" is a clause, e.g. "a rabbit run")

The first analysis is that favored by sense-datum theorists. The problems it faces are well known, and I will not repeat them here. I only want to remark on two points which seem to me to refute it beyond any doubt.

The first point is that the proposed analysis does not capture the difference between merely seeing an F ψ and seeing an F both ψ and ϕ, whatever the additional properties are. Right now I see a can of Coca-Cola sit on my desk. According to the proposed analysis, this means that there is an x which is a Coke can, sits on my desk, and is seen by me. But the can also has the property of having been produced by Amatil Inc from Sydney. If we apply the analysis from right to left, we get that I am seeing a Coke can made by Amatil Inc sit on my desk. This is true on a material reading, but we are interested in the logical form of the intensional phenomenal reading, and such a reading must capture the nature of my experience and nothing more. The statement about Amatil Inc does not. It does not help to suppose that

the individuals we see in phenomenal seeings are mental sense data: whatever the objects, they will have infinitely many properties we can't experience.

This problem could be remedied to some extent by restricting the analysis to properties one can see. One might say:

> I see an F $\psi = \exists x(F(x) \wedge \psi(x) \wedge see(I,x))$, where F and ψ are visually experiencable properties.

This helps, but we can still reasonably expect the account to overgenerate. Note first that we would have to count as visually experiencable all properties which can in principle be visually experienced. Otherwise the analysis would not cover some possible kinds of visual experience. Now suppose (as is plausible) that the property of being very far away from a small object is (in principle) perceptible visually. Arguably everything which actually exists and can be perceived has this property, including mental particulars (if any): everything which exists in a given corner of the universe is very far away from some particle in a different corner of the universe. Given what we have established so far, it follows on the present analysis that whenever I see_{ip} something, I see_{ip} something very far away from a small object. This strikes me as obviously false.

The second problem is that the immaterial individuals sense-datum theorists need to posit in order to account for hallucination clearly can't bear such properties as *being bigger than an elephant*. Yet one can see something being of a given size without there being anything of that size. This requires that the analysis be changed to make use of "primed properties":

> I see_{ip} an F $\psi = \exists x(F'(x) \wedge \psi'(x) \wedge see(I,x))$

I find that I am short of spare concepts to define the required primed properties. For each property that I have a concept for which is such that I could conceivably be said to see_{ip} something bearing it, I can readily imagine having a hallucination of something having that property. In other words, for any given property, either I

cannot make sense of seeing$_{ip}$ something having it, or I can see$_{ip}$ something having it without there being anything which has it. This leaves no room for primed properties.

The second way of cashing out the idea that the relata of phenomenal B-perceivings are individuals is even less appealing. It requires that everyone who sees a pink elephant fly stands in the seeing relation to the very same particular.[5] One can see a pink elephant fly without there being any flying pink elephant or trope of the property of being a flying pink elephant, so the particular in question could not be a flying pink elephant or a trope of being a flying pink elephant. This makes it hard to see what it might be.

Let us now turn to the property account of the objects of phenomenal B-perceivings. On this approach, one might say for example that someone who sees$_{ip}$ a pink elephant fly stands in the seeing relation to the first-order property of being a flying pink elephant.

Problems arise for this view once we consider slightly more complex examples:

(4.26) I see a large number of bright particles floating in front of me.

(4.27) I see two entities floating in front of me.

Case 4.26 is problematic if we read it as ascribing a state of seeing an indeterminate number of particles (as in the speckled hen case). It is arguably possible to experience (hence see$_{p}$) an indeterminate number of objects. How could an object property characterize one's state in such cases? To say that what one sees is the property of being one of an indeterminate number of objects would seem to misconstrue the character of the experience: in some cases at least, when one sees an indeterminate number of objects, there is not a single object that one determinately sees.

[5]This constraint could perhaps be relaxed, but this would break some of the inferences we discussed in the preceding section. There is also no comfortable resting place between this and saying that every experiential event involves a relation to a different particular, which is nearly unintelligible (what about overlapping events?).

Case 4.27 is even more problematic. If what I see in this case were a property, it would have to be the property of being two entities floating in front of me. But no individual could possibly have the property of being two entities. My experience would therefore be necessarily false. It clearly is not, so what I see is not a property.

The preceding problems could conceivably be addressed by taking the relevant properties to be properties of states of affairs instead of properties of individuals. The relevant properties might be said to be *involving a large number of bright particles floating in front of me* and *involving two entities floating at such and such locations.* But properties of states of affairs are arguably the same as properties of the world, and such properties count as propositions for our purposes. This is not an entirely ad hoc maneuver, because properties of the world are naturally described as sets of possible worlds, or using predicate logic (our criterion of propositionhood). Also, some theorists explicitly take propositions to be world properties (c.f. Soames 2005).

Another problem with the property approach is that it does not seem to do justice to the inferential role experiences play. There is a range of views about the extent and manner of the justification provided for everyday beliefs about the external world by experience, but nearly everyone seems to agree that experiences can in some sense be "endorsed" or "taken at face value", and it seems that the judgments arrived at by endorsing our experiences often have existential import. If I take my current visual experience at face value, for instance, I will judge that I am in presence of a bright surface. It is hard to see how I could arrive at this judgment through mere endorsement if my experience related me to a mere uninstantiated property. Why would I arrive at an existential judgment rather than a universal one? Or why not conclude that there are five objects instantiating the property?

Of course, a causal explanation of the quantified judgments we arrive at through experience might still be possible. But a *rational* explanation of a perceptual judgment to the effect that there are two elephants in front of oneself when one merely

experiences the property *being an elephant* does not seem possible. Arguably, for such an explanation to be possible, there must be logical implication from the content of the experience to the content of the judgment, or one must have some background beliefs which justify the judgment in light of the fact that one has the experience. The first condition is not meant by experiences which have properties of individuals as contents. The second condition would typically not be meant either: who believes that experiences of the property *being an elephant* normally occur in presence of *two* elephants? On the other hand, if my experience relates me to a proposition, the content of my belief when I endorse my experience can simply be the content of my experience. My belief that there are two elephants in front of me can be justified by my experience of the fact that there are two elephants in front of me.[6]

A more theoretical difficulty with properties is that it is hard to see how such an account could lend itself to a compositional explanation of the variety of experiences we can undergo. As noted in chapter 2, the large diversity of experiences we can undergo is likely generated by some combination mechanism. On the propositional approach, it is reasonably clear how this could work at an abstract level. For the sake of illustration, we can imagine a kind of language of experience along the lines of Fodor's Language of Thought. What one experiences would be a proposition corresponding to a sentence in this language. The contents of phenomenal sentences would be determined through recursive rules operating on their constituents. If what we experience were properties, on the other hand, it is hard to see how this kind of compositional system might work. An explanation based on this account of the contents of experience would have to begin by articulating the internal structure of properties, which is far more elusive than the internal structure of propositions.

One might hope to avoid some of the preceding difficulties by positing that what we experience are sets of properties rather than properties. One problem with

[6]One might think that a virtualist cannot consistently talk of experiencing facts. I will address this objection in section 7.2.

this approach parallels the many-property problem Jackson (1975) raises against adverbialism. Take for example the state of seeing$_{ip}$ a yellow circle to the left of a green square. If we cash this out as seeing$_{p}$ the set {yellow, circle, left of, green, square}, we lose all distinctions between different ways of seeing$_{p}$ these properties, for example, between the aforementioned experience and seeing$_{ip}$ a yellow square to the left of a green circle. The same goes if we take the set {yellow circle, left of, green square}. In this case, we conflate the two different positions the two objects experienced could be in. To capture the precise character of one's experience, we must invoke a set containing a single all-encompassing property of *being a yellow circle on the left of a green square*, which takes us back to the single property approach, but with a superfluous formal complexity.

Most of the preceding problems could easily be avoided by an account which takes the objects of experiences to be more exotic entities than we have considered so far. For example, perhaps higher-order properties have the requisite fineness of grain to avoid some of the problems I raised for the first-order property approach. If not, sets of propositions will probably deal with all the problems. My only qualm with these accounts is that they seem phenomenologically inaccurate. It does not seem far fetched to say that we experience propositions, provided that we count state of affairs-like entities under this heading. As we will see in chapter 7, disjunctivists rest their position almost entirely on the putative phenomenological datum that experience acquaints us with simple, first-order states of affairs. I will argue in chapter 7 that the propositional approach, broadly construed, can largely accommodate this observation. On the other hand, it seems prima facie phenomenologically inaccurate to say that our basic experiences are experiences of properties of properties, set of propositions, or more complex abstract objects. Since the propositional approach also has the advantage of simplicity, I think that the problems we uncovered for the simple accounts discussed so far strongly militate for this approach.

The conclusion of the preceding section was that phenomenal perceivings are

states of standing in modality-specific relations to certain entities. In this section, I have argued that, in the case of phenomenal B-perceivings, the relevant entities are propositions. When I say that I see_{ip} an elephant flying, for example, I am expressing the proposition that I stand in the seeing_{p} relation to the proposition that an elephant is flying, or:

$$see_p(I, [\exists x(elephant(x) \wedge flying(x))]),$$

where [S] denotes the proposition symbolized by formula S (as in Bealer 1979). This completes my case for sensory propositionalism.

4.5 The virtual character of phenomenal perceivings

So far I have focused exclusively on the case for sensory propositionalism, the first component of sensory virtualism. To establish sensory virtualism, we also need to show that the relations which are constitutive of phenomenal B-perceivings are virtual.

For this, we need only refer back to section 3.2.2, where I argued that phenomenal perceivings are states one can be in independently of whether what one perceives obtains or not. I gave realistic (if not real) examples which illustrate this. These examples show that the seeing_{p}, $\text{hearing}_{\text{p}}$, ecetera relations are virtual. I will discuss the virtual character of phenomenal perceivings at greater length in chapters 7 and 8 in response to the case for disjunctivism.

This completes my case for sensory virtualism. In the next section I will discuss the semantics of D-perceivings. While not essential to the virtualist position, the points in this section throw useful light on a number of issues I have so far left unaddressed and provide additional motivations for sensory propositionalism.

4.6 The semantics of D-perceiving ascriptions

When I introduced intensional readings of perceptual ascriptions in section 3.1, I defined them as readings which do not have a certain logical form. I did not offer a positive characterization of their logical forms. The account of phenomenal B-perceivings I gave in the preceding section answers this question as far as ascriptions of phenomenal B-perceivings go. In this section I will try to shed some light on the case of phenomenal D-perceivings. I will not talk about non-phenomenal D-perceiving ascriptions—there are too many different kinds, and they are not directly relevant to our purposes here. I will conclude the section by drawing some conclusions regarding the relation between phenomenal D-perceivings and phenomenal B-perceivings.

In section 3.3, we saw that there are plausibly two types of ascription of phenomenal D-perceivings: "NP ϕ-s DP" and "NP ϕ-s N" ascriptions, where "DP" stands for a determiner phrase and "N" stands for a noun. There are other kinds of ascription of D-perceivings (e.g. "that"-clause ascriptions), but it is plausible that they do not lend themselves to phenomenal readings, so I will set them aside. Assuming that perceptual verbs denote relations in ascriptions of phenomenal D-perceivings just like in all other perceptual ascriptions discussed so far, the question is what one is related to in phenomenal D-perceivings. For example, assuming that "sees" denotes a relation in "Alice sees_{ip} a rabbit", the question is what Alice has to be related to in order to see_{ip} a rabbit. I will discuss "NP ϕ-s DP" ascriptions first, then I will talk briefly about the general case of "NP ϕ-s NP" ascriptions.

The potential objects of the phenomenal D-perceivings ascribed with DP complements are the same as in the case of phenomenal B-perceivings. Individuals and properties of individuals raise the same problems as before with minor adaptations—individuals yield material readings, and properties of individuals are too coarse-grained. Propositions are excluded because they are the objects of B-perceivings and B-perceivings are by definition not D-perceivings. We must therefore consider

more complex or abstract entities.

A good candidate kind of entity is *generalized quantifiers*. In logic, generalized quantifiers are characterized as sets of sets of individuals (see Barwise and Cooper 1981, Westerståhl 2008). That is to say that they are second order properties. To keep things simple, we can restrict our attention to generalized quantifiers which can be described with lamba formulas of the form

$$\lambda Q(\psi(P(x) \wedge Q(x))),$$

where ψ stands for a regular quantifier ("some", "every") and the "P(x)" conjunct is optional. For example, the entity one is related to when one sees_{ip} a rabbit could be the property $\lambda Q(\exists x(rabbit(x) \wedge Q(x)))$. In words, this is the property of being a property had by a rabbit.

Montague (1973) treats transitive intensional verbs with DP complements as denoting relations to generalizer quantifiers. His approach is now widely accepted (with some variations) among semanticists (Dowty et al 1981; Forbes 2008; Westerståhl 2008).[7] One of its main virtues is that it meshes well with the fact that "NP V DP" ascriptions allow both intensional and material readings. More specifically, the difference between intensional and material readings turns out to be simply a difference in the order of predication on this view as far as DP ascriptions go. Here I am going beyond the standard treatment, but this claim is not hard to verify. Consider the property of being seen by me, $\lambda x(see(I,x))$. If we predicate this property of $\lambda Q(\exists x(rabbit(x) \wedge Q(x)))$, the resulting proposition reduces to

$$see(I, \lambda Q(\exists x(rabbit(x) \wedge Q(x)))),$$

which is just the proposition expressed by "I see a rabbit" on the intensional reading of this statement according to the Montague-style account. Reversing the order of predication and predicating $\lambda Q(\exists x(rabbit(x) \wedge Q(x)))$ of $\lambda x(see(I,x))$, we get the proposition

[7] Interestingly, Montague (1973) does not extend this account to perceptual verbs, because he does not treat them as intensional verbs. The reasons for this are not clear.

$$\exists x(rabbit(x) \wedge see(I,x)),$$

which is just the proposition expressed by "I see a rabbit" on the material reading of this statement. On Montague's treatment of DP ascriptions, the ambiguity between intensional and material readings corresponds to an ambiguity in the order of predication.

Montague's treatment of DP ascriptions can be extended to all NP ascriptions. In section 3.1 I noted that a statement such as "Alice sees Bob" might seem to have no material reading if it is always read as expressing the proposition $see(Alice, Bob)$. But in modern semantics, names do not denote individuals. Rather, they denote generalized quantifiers—all NPs denote generalized quantifiers (Barwise and Cooper 1981;Westerståhl 2008; Ruys and Winter 2001). The quantifier denoted by "Alice", for example, can be represented as $\lambda Q(\exists x(Alice*(x) \wedge Q(x)))$. This account is primarily motivated by its ability to explain how simple nouns combine with quantified noun phrases, for example, in 4.28 (ibid.).

(4.28) Alice and her lawyers wanted to call.

If "Alice" and "her lawyers" both denote second-order properties (sets of properties), the denotation of the sentence's subject can be derived through set operations. In this case it would be the set of properties had by Alice and all her lawyers (if any).

On the present account, we would expect "Alice sees_{p} Bob" to have at least two readings:

(4.29) $\exists x(Alice*(x) \wedge see_p(x, \lambda Q(\exists y(Bob*(y) \wedge Q(y)))))$

(4.30) $\exists x(Alice*(x) \wedge \exists y(Bob*(y) \wedge see_p(x,y)))$

4.29 is an intensional reading, while 4.30 is a material reading. 4.29 illustrates the logical form a pure phenomenal ascription of the form "NP sees N" would have, however counterintuitive this reading might be (I am not convinced that such states of affairs really are possible; see 3.1 and 7.2).

Extending our account of DP ascriptions to all NP ascriptions has the side benefit of legitimizing my decision to call intensional perceivings "intensional". As I acknowledged in section 3.1, my choice of terminology might have seemed inappropriate in light of the common belief that a statement such as "Alice sees Bob" should always be read as expressing the proposition $see(Alice, Bob)$. But it turns out that this statement has a material reading (4.30), and that this reading is by far the most natural. My material / intensional distinction turns out to be in line with prior usage.

One significant virtue of the preceding account of phenomenal D-perceivings is that it yields a simple explanation of how phenomenal B-perceivings realize phenomenal D-perceivings (recall that phenomenal D-perceivings are derivative). It seems plausible that a D-perceiving which consists in standing in a relation R to a generalized quantifier ψ is realized by a B-perceiving which consists in standing in R to a proposition P when P *contains* ψ. To give a full account of when a proposition contains a quantifier would require an account of the structure of propositions, but it is plausible on a first pass that P contains ψ just in case P entails that something (a property) instantiates ψ. Take for example the quantifier designated by "a rabbit" ($\lambda Q(\exists x(rabbit(x) \wedge Q(x)))$). The proposition expressed by "a rabbit runs" entails that some property instantiates the quantifier designated by "a rabbit", namely, the property of running. So $\text{seeing}_{\text{ip}}$ a rabbit run is one way (one realizer) of $\text{seeing}_{\text{ip}}$ a rabbit.

There is of course a risk that this account overgenerates. The proposition expressed by "a rabbit runs" entails not only that something has the quantifier designated by "a rabbit", but also that something has the quantifiers designated by "a proposition", "a number", "a set", etc. It entails every necessary truth, hence realizes any state of seeing ψ, where ψ is a quantifier occurring in a necessary truth of the form "ψ is F". So when you see a rabbit, you also see a set, a proposition, a number, etc. I don't find this to be obviously unacceptable. I suspect that this

merely shows that there is not much it is like to see any of these things. If seeing any of these things amounts to nothing more than seeing the unique unstructured necessary proposition or state of affairs, we should not expect that there is much associated phenomenology. It does not seem far-fetched to suppose that all phenomenal experiences involve in the background an experience of The True.

In any case, it should be possible to impose more constraints on a B-perceiving realizing a D-perceiving if required. To do this, we would need a theory of the structure of the propositions we experience. I cannot provide such a theory here, but I don't see any principled difficulty. Either the propositions we experience have structure or they do not. If they do, we should be able to uncover it. If they do not, then the preceding account should be sufficient.

This account of phenomenal D-perceivings and their grounds seems not only well motivated, its coherence with our propositionalist account of B-perceivings arguably lends further support to the latter. However, one obvious objection need be addressed. Earlier I rejected higher-order properties as objects of B-perceivings on phenomenological grounds (among others). I am now saying that phenomenal D-perceivings are relations to higher-order properties. Should not the phenomenological point carry over to D-perceivings?

My point against higher-order properties only applies to basic phenomenal states. The point is that when one introspects the most basic phenomenal states one can undergo, it never occurs to one that they might be relations to higher-order properties or more complex or abstract entities. If they seem relational, they seem at most to relate us to proposition-like things with fairly simple first-order structure. The point does not extend to derivative phenomenal states (D-perceivings) because these are not directly introspectible. When you notice that you see_{ip} something, you notice this by noticing that you see_{ip} something having certain features. Then you can choose to describe your state more abstractly as a mere $\text{seeing}_{\text{ip}}$ of "something". Since derivative phenomenal states are essentially abstractions of basic phenome-

nal states, it befits them to have a more abstract structure.

4.7 Summary

In this chapter I presented a prima facie case for sensory virtualism based on PCSE+, the view that the basic phenomenal states are phenomenal B-perceivings. I began with a simple argument to the effect that sensory propositionalism, the first tenet of sensory virtualism, offers the most plausible account of the semantics of ascriptions of phenomenal perceivings which appear to quantify over propositions. I then elaborated a more complex argument which proceeds by elimination of the two main alternatives to sensory propositionalism: the view that phenomenal perceivings are monadic states, and the view that phenomenal B-perceivings are relations to non-propositions. Given that phenomenal perceivings are independent of one's environment (as argued in section 3.2.2), sensory propositionalism implies sensory virtualism. I concluded the chapter with a brief discussion of the semantics of D-perceivings ascriptions and the relation between D-perceivings and B-perceivings.

Chapter 5

Pure virtualism

In the previous chapter, I presented a case for sensory virtualism (SV). The virtualist theory I want to defend is strictly stronger than SV:

(Pure) virtualism There is a relation R such that: 1) For any basic phenomenal state *s*, there is a proposition P such that *s* = standing in R to P; 2) R is virtual with respect to basic phenomenal states.

For the duration of this chapter, I will refer to this theory as "pure virtualism" in order to avoid any confusion with the weaker SV. I will now argue for pure virtualism based on SV. I will argue that SV is not a good resting point: if one is attracted to SV, one should go all the way and endorse pure virtualism. I will begin by delineating the options we are left with given SV and the considerations which militate for it. Then I will discuss two problems one has to face if one wants to hold to SV while rejecting the pure virtualist view. I will conclude with responses to some arguments which might seem to favor a weaker position than pure virtualism.

5.1 Beyond the senses

SV applies only to experiences which belong to certain sensory modalities. On the face of it, it does not apply to cognitive and emotional experiences, which are not

readily associated with sensory modalities. But it seems highly plausible that all experiences are similar in nature: if sensory experiences consist in standing in virtual relations to propositions, it seems probable that all experiences have this structure as well. The question is how we should generalize from SV to all experiences. We have two options: we can endorse pure virtualism, or we can find some way of extending SV to cognitive and emotional experiences. I will say more about each of these options before arguing for the first.

5.1.1 More on pure virtualism

Pure virtualism does not say anything about the nature of the relation R it posits. As such, the theory is incomplete. And it is hard to see how it could be true without having some idea of what relation R is. Fortunately, we now have all the material we need to specify it.

The relation which I believe makes pure virtualism true is the introspectively salient relation which is part of seeing$_p$, hearing$_p$, smelling$_p$, feeling$_p$, etc. I will refer to this relation as *the experiencing relation*. The experiencing relation is part of seeing$_p$ and other acts of phenomenal perception in that seeing$_p$ *x*, hearing$_p$ *x*, ecetera all consist at least in part in experiencing *x*.

The experiencing relation The introspectively salient relation which is at least partly constitutive of seeing$_p$, hearing$_p$, smelling$_p$, feeling$_p$, etc.

That there is an experiencing relation should be clear from the fact that we systematically describe all phenomenal perceivings as cases of experiencing something, in the technical sense of "experiencing": when one sees$_p$ stars, one experiences stars; when one smells$_p$ a rubber smell, one experiences a rubber smell; when one feels$_p$ a sensation, one experiences a sensation; and so on. I discussed this matter at some length in chapter 3. It is also hard to make sense of the fact that everyone agrees that all phenomenal perceivings have something in common (that all are phenom-

enal states) without positing some property which is constitutive of all of them.[1] Given the relational structure of phenomenal perceivings, it seems quite plausible that the property they have in common is a relation.

To say that there is an experiencing relation as defined above is not yet to say that pure virtualism is true: it could be that experiences involve more than standing in the experiencing relation to certain contents, but also (for example) further elements corresponding to their specific modalities. However, the experiencing relation is a good candidate to play the role of R in the pure virtualist theory. I will refer to the claim that pure virtualism is satisfied by the experiencing relation as *virtualism+*.

Virtualism+ (Pure) virtualism is true, and relation R is the experiencing relation.

One potential difficulty with virtualism+ is that it is not immediately obvious that it is compatible with the considerations which led us to SV. For we found that every basic phenomenal state in a given modality ϕ is a state of standing in the ϕ-ing relation to a proposition. Can a phenomenal state be both a state of standing in the seeing relation to P and a state of standing in the experiencing relation to P?

Yes, provided that we can in principle give a pleonastic analysis of the seeing relation. To illustrate what I mean, take statements 5.1 and 5.2.

(5.1) Alice ran a marathon.

(5.2) Bob is circling the roundabout.

There is arguably nothing more to running a marathon than participating (perhaps properly) in a marathon, because a marathon is an event in which one runs. Similarly, there is nothing more to circling a roundabout than following the contour of a roundabout, because roundabouts are (at least roughly) circular. On the virtualist+ view, the experiencing relation is to seeing$_{ip}$ something red what the act of following

[1] Disjunctivists sometimes say that veridical and non-veridical experiences have *nothing* in common. This is a logical falsehood. What disjunctivists really mean is that veridical and non-veridical experiences do not share properties of a certain kind (most likely that they do not share phenomenal characters).

the contour is to circling a roundabout: there is nothing more to seeing$_{ip}$ something red than experiencing$_{i}$ something red, because to see$_{p}$ is just to experience an entity of the visual kind, and what "something red" denotes in the "see$_{ip}$ ___" context is an entity of that kind.

What is a visual (auditory, tactile, etc.) proposition or quantifier? On a first pass, it seems that a visual proposition is one that involves color properties and no other properties but modality-neutral properties (e.g. spatial and temporal properties). Aural entities might be characterized by their involvement of sound properties (e.g. pitch). It is less clear what the properties characteristic of touch and other modalities might be. But the matter is of little importance from the standpoint of the pure virtualist theory, because modalities are not part of the theory. For the pure virtualist, modalities do not cut nature at her joints. It might be hard to make explicit exactly what the folk mean by "visual", "auditory", ecetera, but this is not a task that we need to achieve in order to develop a complete virtualist account of phenomenal character.

5.1.2 Impure virtualism

Pure virtualism(+) is one way of generalizing from SV to all phenomenal states. But this theory might well seem too strong. As noted at the end of chapter 1, some representationalists hold that an experience's modality affects its phenomenal character independently of its content. Visually experiencing a shape and tactilely experiencing the same shape, for example, might be thought to be two phenomenal states with the same content but distinct phenomenal characters. Pure virtualism does not allow this. So before we adopt pure virtualism, we need to consider a weaker generalization of SV which gives some role to modalities or modality-like factors in the determination of phenomenal characters. The view Chalmers (2004) calls *impure representationalism* is the representational theory which best matches this job description. We can call its counterpart in the virtualist framework *impure*

virtualism.

Impure virtualism (IV) For any manner of representation M, there is a relation R such that: 1) For any basic phenomenal state *s* which has M there is some proposition P such that *s* = standing in R to P; 2) R is virtual with respect to basic phenomenal states.

We can leave the question of what a manner of representation is largely open for now. For the moment, let us assume only that every phenomenal state has a manner of representation. Without this assumption, IV would not cover all phenomenal states, so it would not really be a competitor to pure virtualism *qua* generalization of sensory virtualism. We can for the time being take manners of representation to correspond to sensory modalities and other kinds of experience we have found it useful to distinguish so far, e.g. cognitive and emotional experiences.

IV, like SV, is strictly weaker than pure virtualism. It is not an opposing view but an alternative stopping point. We can call the position according to which IV is correct but pure virtualism is not "IV-".

IV- IV is true, but pure virtualism is false.

There is good reason to think that some generalization of SV is true. Allowing for sufficiently many interpretations of the notion of a manner of representation, this means that either pure virtualism is true or IV- is true. In the rest of this chapter I will argue for pure virtualism principally by arguing against IV-.

5.2 The case for pure virtualism

Ultimately, the primary reason why I favor virtualism+ over IV- is that the former is simpler, clearer, and has considerably more explanatory power. But virtualism+ also seems more phenomenologically accurate to me. I want to begin by trying to bring this out.

Imagine that you are looking at a Coke can on your desk. When you see_{ip} a red can on your desk, you experience a certain scene: you stand in the experiencing relation to a state of affairs such as $\exists x\, redCan(x)$. You can easily recognize something that is going on which is also going on when you hear_{p}, smell_{p}, taste_{p} or feel_{p} things: in each of these cases, as in your Coke can experience, you stand in the experiencing relation to something. Now try to find some aspect of the phenomenology of your state of seeing_{p} a red can which goes beyond the fact that you are experiencing a red can. Take a moment to try this.

One must not conclude that there is a phenomenal visualness to the seeing_{p} state above and beyond its being an experience of a red can on the ground that it is a specifically visual experience. For being a visual experience might just be a matter of being an experience with color content (as the pure virtualist posits). Personally, however carefully I introspect, I am unable to discern anything in $\text{seeing}_{\text{ip}}$ a red can which is not also part of what it is to experience a red can (to stand in the experiencing relation to a state of affairs involving a red can).

In the rest of this section I will discuss two difficulties with IV- which militate for pure virtualism independently of the preceding phenomenological considerations. These problems come out more clearly if we assume that IV- implies that there can be phenomenal states which have the same content but differ in phenomenal character. While IV- does not strictly speaking imply this, a proponent of IV- has to accept this. The reason is that there would be no clear counterexample to pure virtualism otherwise.[2] As we will see in section 5.3, the case for IV- rests entirely on putative differences in phenomenology between experiences in different sensory modalities which do not correspond to differences in content. If there were no such examples, pure virtualism would seem to be by far the best theory given its simplicity. I will therefore assume that IV- commits one to there being distinct phenomenal states with the same contents. Another assumption which I feel licensed

[2]That is, no clear counterexamples which leave IV intact.

to make at this stage is that we should reject IV- unless there is an interpretation of IV which makes it a) plausible and b) at least as strong as sensory virtualism. There is good reason to think that sensory virtualism is merely a special case of the true general account of the structure of phenomenal states. If no interpretation of IV makes sensory virtualism a special case of IV, then there is good reason to think that the best characterization of the structure of phenomenal states is that provided by pure virtualism.

5.2.1 The problem of multimodal experiences

Phenomenology is generally unified. Right now, for instance, I seem to have a unified visual experience which brings together or subsumes all my other visual experiences. The same relation of subsumption holds between experiences across sensory modalities: I see the keys on my keyboard, I hear the noise they make as I hit them, and I also experience the noise together with the keys. I seem to have *multimodal experiences*: experiences which subsume simpler experiences in different sensory modalities.

The notion of experience subsumption is hard to analyze, but one thing that is fairly clear is that it requires entailment or necessitation: if e_1 subsumes e_2, then e_1 necessitates e_2. Intuitively, if an experience e_1 subsumes an experience e_2, e_2 is a phenomenal part of e_1. Since all phenomenal aspects of e_1 and e_2 are essential to them, it would seem to follow that e_2 is an essential part of e_1 if e_1 subsumes e_2. It would therefore seem to follow that e_1 necessitates e_2 if it subsumes it.

For present purposes, I will assume that there is nothing more to subsumption than necessitation: e_1 subsumes e_2 just in case e_1 necessitates e_2. See Bayne & Chalmers (2003) for more on subsumption and a defense of this equivalence. For our purposes, we need not think of the equivalence between subsumption and necessitation as a substantive thesis; we can think of it simply as defining one kind of subsumption or unity. The multimodal experiences I am interested in here are sim-

ply the phenomenal states which necessitate phenomenal states in different sensory modalities.

In some cases, it might seem that multimodal experiences are nothing more than conjunctions of experiences in different modalities. My unified experience of the noise and the keys of my keyboard, for example, might be thought to be nothing more than the state of experiencing the keys at the same time as experiencing the noise they emit. However, other multimodal experiences are harder to reconstruct in this fashion. For example, it seems that I can experience a noise *as* coming from a dark shape in my environment, where I am experiencing the noise aurally and the shape visually. Call this experience e.

On the pure virtualist view, it is relatively easy to explain how an experience like e can subsume experiences in different modalities. Let us say for example that the content of e is 5.3.

(5.3) $\exists x \exists y(noise(x) \wedge darkShape(y) \wedge caused(y,x))$

There are two kinds of visual and auditory experience e might reasonably be said to subsume. First, e plausibly subsumes two D-perceivings: seeing$_{ip}$ a dark shape and hearing$_{ip}$ a noise. Second, we could also say that e subsumes the B-perceivings seeing$_{ip}$ something be a dark shape and hearing$_{ip}$ something be a noise.

According to the account of phenomenal D-perceivings I gave in section 4.6, the subsumed experiences' contents in the first case would be quantifiers 5.4 and 5.5.

(5.4) $\lambda Q(\exists x(noise(x) \wedge Q(x)))$

(5.5) $\lambda Q(\exists x(darkShape(x) \wedge Q(x)))$

I suggested that a phenomenal B-perceiving involving a virtual relation R realizes (hence necessitates) a phenomenal D-perceiving also involving R when the former's content entails that the latter's content is instantiated. If pure virtualism is correct, the present experiences all involve the same virtual relation. The condition of entailment also obtains between their contents: 5.3 entails that 5.4 is instantiated by the

property of being caused by a dark shape and that 5.5 is instantiated by the property of causing a noise. Our account of realization therefore provides an account of e's unity if the experiences it subsumes are D-perceivings.

The components of e could also have full propositional contents (they could be B-perceivings):

(5.6) $\exists x\, darkShape(x)$

(5.7) $\exists x\, noise(x)$

In this case the account of realization given in chapter 4 does not apply, because it only covers the cases of derivative phenomenal states. However, basic phenomenal states can also necessitate one another, and our current account is naturally extended to this case. To generalize the account, we can say that a phenomenal state s_1 necessitates a phenomenal state s_2 when that the content of s_1 obtains entails that the content of s_2 obtains. An immediate implication of this is that a basic phenomenal state s_1 necessitates a basic phenomenal state s_2 when the former's content entails the latter's content. If this account is acceptable, the pure virtualist has a straightforward explanation of how an experience of content 5.3 can unify experiences of contents 5.6 and 5.7, because proposition 5.3 entails propositions 5.6 and 5.7.

Of course, the same overgeneralization worries we discussed in chapter 4 arise. But the same response applies. The account has the consequence that we experience every necessary truth all the time, but one way to see this is as proof that the phenomenology of an experience of a necessary truth is always the same, and that there is not much to it.

If this account turns out to be too liberal, an alternative is to say that a basic phenomenal state s_1 necessitates a basic phenomenal state s_2 when the former's content *contains* the latter's content, where containment is something more demanding than entailment. As noted in chapter 4, spelling this out would require giving an account of proposition structure, but we may expect something along these lines to be

correct if the simple entailment account proves untenable.

Now, this purely content-oriented way of deriving simple experiences from more complex ones works because all experiences are constituted by the same virtual relation on the pure virtualist view. Within IV-, we need to add an explanation of how a subsuming experience's manner of representation relates to the manners of representation of the experiences it subsumes.

The main difficulty here is to give a plausible account of the virtual relations which are constitutive of multimodal experiences. It is unclear what the virtual relation constitutive of *e* could be. It is not just a visual experience or just an auditory experience. Should we say that it is both auditory and visual? There is a sense in which it is, but the visual and auditory relations are distinct on IV-. We must either choose which is constitutive of *e* or introduce a new relation for visual-and-auditory experiences (and all other combinations of modalities). It seems prima facie implausible that *e* is constituted by the auditory or visual relation (but not the other), so we need to give an account of the combined visual-and-auditory relation which constitutes it.

Various accounts of the visual-and-auditory virtual relation are possible. For example:

(5.8) $\lambda x \lambda P$*(x represents proposition P in part visually and in part aurally)*

(5.9) $\lambda x \lambda P$*(x represents proposition P both visually and aurally)*

5.8 suffers from what we might call the *wholeness problem*: it does not seem to do justice to the fact that, in *e*, one does not merely experience a part of P visually while experiencing another part of it aurally: one also experiences the whole of P as well, including the relation of causation between the noise and the shape (which is not clearly accounted for on 5.8). In response to this, one might suggest 5.10.

(5.10) $\lambda x \lambda P$*(x represents proposition P in part visually and in part aurally, and in whole neutrally)*

But 5.10 suffers from the *duplication problem*: it seems to imply that one has two phenomenally distinct experiences of the noise involved in P (say): one auditory experience of the noise, and some other "neutral" experience of the noise. Whatever the neutral manner of representation is supposed to be, it should make for different phenomenal states than the auditory manner. For the virtual relation associated with the neutral manner would satisfy pure virtualism if it yielded the same phenomenal states as all other manners of representation.[3] But I don't have two distinct experiences of exactly the same noise when undergoing an experience like e, so this seems wrong. 5.9 suffers from the same problem: it seems to imply that in e one experiences the whole of P both visually and aurally, while in fact one experiences it only once, but in a mixed way.

The preceding accounts of multimodal relations all face an additional difficulty. Unless different manners of representation (or constitutive virtual relations) can be associated with the same content, IV- is implausible. Let us say for the purposes of illustration that simple shape contents yield different experiences when they are represented in touch than when they are represented in vision. Then the two following multimodal phenomenal states should be possible:

> e_1: A state in which one visually experiences a square, tactilely experiences a circle, and experiences the square as being on top of the circle.
>
> e_2: A state in which one tactilely experiences a square, visually experiences a circle, and experiences the square as being on top of the circle.

Again for the purposes of illustration, we can suppose that these states have the content $\exists x \exists y(circle(x) \wedge square(y) \wedge onTopOf(y,x))$. Call this proposition P. e_1 and e_2 are both at once visual and tactile. On any of the preceding accounts, both would be states of standing in the visual-and-tactile virtual relation to P. That is to say that they would be the same state. But e_1 and e_2 are clearly distinct.

[3] Of course, one might say that the neutral manner happens to yield the same phenomenal state as the auditory manner in the case of simple noise content. But the two manners of representation have to differ for some contents, and we could modify the example accordingly.

This problem calls for an account of the virtual relations involved in multimodal experiences which individuates them more finely. One might think that e_1 is a state in which one is related to P through a relation such as 5.11.

(5.11) *λxλP(x represents proposition P visually as far as squares go and tactilely as far as circles go)*

The general idea is to specify what kinds of entity are experienced in what modalities as part of the virtual relation. But one can presumably experience two entities of the same kind in different modalities as part of the same experience (if IV- is correct). For example, a phenomenal state such as the following should be possible if IV- is correct:

> e_3: A state in which one visually experiences a square, tactilely experiences a square, and experiences the first square as being on top of the second square.

5.11 does not work with experiences like e_3. We would have to say that e_3 involves a virtual relation such as this one:

(5.12) *λxλP(x represents proposition P visually with regard to the square at location L1 and tactilely with regard to the circle at location L2)*

At this stage we have built all the content of the experience into the virtual relation. This strikes me as highly ad hoc.

Another problem with both 5.11 and 5.12 is that they seem inconsistent with sensory virtualism. It is reasonable to suppose that if multimodal experiences involve complex, shape- and location-involving virtual relations like 5.11 and 5.12, so do unimodal experiences. But this would seem to contradict sensory virtualism, which asserts that all experiences in the same modality are constituted by the same virtual relation.

Whether 5.11 and 5.12 contradict sensory virtualism or not, they seem to require an interpretation of the manners of representation referred to in IV which makes the

latter weaker than sensory virtualism in not assigning the same virtual relations to experiences in the same modalities. As I said at the beginning of this section, there is good reason to think that there is some true generalization of sensory virtualism which captures all its force. If the only way of making IV work is to make it weaker than sensory virtualism in some respect by allowing experience-specific virtual relations, the only suitable generalization of sensory virtualism is that provided by pure virtualism.

Note also that 5.11 and 5.12 face the same dilemma between the wholeness problem and the duplication problem simpler approaches face. Relations 5.11 and 5.12 do not capture the fact that the shapes experienced in states e_1 and e_3 are experienced as related to each other. To remedy this, we need additional clauses to the effect that one represents the whole of the contents of e_1 and e_3 in a "neutral" manner. But then it follows that one experiences the relevant shapes twice. This dilemma is perhaps the most difficult aspect of the problem of multimodal experiences for impure virtualism.

Aside from the problem of specifying the virtual relations constitutive of multimodal experiences, another difficulty multimodal experiences like e generate for IV- is that of explaining how they necessitate their components. The difficulty lies in explaining how one's being in e could entail that one sees_{ip} a dark shape without also entailing that one sees_{ip} a noise, or how it could entail that one hears_{ip} a noise without also entailing that one hears_{ip} a dark shape. To show that e necessitates one of its component experiences e_{i}, one would have to show that α's being in e, possibly together with additional premises which are necessarily true, entails α's being in e_{i}. In order to show this, one would have to produce a suitable valid argument. Now consider statements 1-3 below, where R_{e} is the virtual relation involved in e. One would have to sketch a valid argument from 1 and possibly some necessary truths to 2, but which does not also yield 3.

1. α stands in R_{e} to $\exists x \exists y(noise(x) \wedge darkShape(y) \wedge caused(y,x))$.

2. α stands in the auditory virtual relation to $\exists x\, noise(x)$.

3. α stands in the auditory virtual relation to $\exists x\, darkShape(x)$.

The same applies to the visual experience subsumed by e: one would have to show that being in e entails standing in the visual virtual relation to $\exists x\, darkShape(x)$ but not $\exists x\, noise(x)$, even though both contents are related to the content of e in exactly the same way. It is very hard to see how this could be achieved; parity of reasoning suggests that it cannot be that 1 entails 2 but not 3.

Here it is important to keep in mind that a proponent of IV- is committed to there being some contents which can be represented in more than one manner of representation. Given that this is so, it would not do to describe R_e as the visual-and-auditory relation and add the following as a (putatively necessary) premise:

> If α stands in the ϕ-and-ψ virtual relation to a content which contains a content C compatible with modality ϕ, then α stands in the ϕ virtual relation to C.

Suppose for the sake of illustration that IV- allows shapes to be represented both aurally and visually. The above premise combined with 1 would entail both 2 and 3. Of course, this particular counterexample can be avoided by stipulating that shapes cannot be represented aurally (and they plausibly cannot). But we could construct a counterexample that fits whatever contents make exception to pure virtualism according to a proponent of IV-. IV- can explain subsumption using the above principle only if every content is always represented in every manner it can be represented in when it is represented at all. This is highly dubious on the assumption that manners of representation affect the phenomenology of experience independently of content as proponents of IV- claim. Ultimately, the only way of explaining the entailment of subsumed experiences by subsuming experiences would seem to be to ascribe to unifying experiences highly complex, content-embedding virtual relations such that we ended up with for e_1 and e_3 above.

At this point it is tempting to reject the assumption we have been making that multimodal experiences like e have the same kind of relational structure as the unimodal experiences they subsume. For example, one might postulate that e is not a state of standing in a virtual relation to a content, but some kind of compound of visual and auditory experiences.

But multimodal experiences like e, e_1, e_2 and e_3 are not mere conjunctions of unimodal experiences. In particular, e is not merely the state of hearing$_{\text{ip}}$ a noise and seeing$_{\text{ip}}$ a dark shape, but a state in which one experiences a dark shape causing a noise. Neither can we obtain it from simpler experiences stuck together with the irreducible phenomenal glue Dainton (2000) calls "co-consciousness": how and why would co-consciousness between an experience of a noise and an experience of a shape result in an experience of the noise as caused by the shape? The problem is that a variety of relations can replace causation in experiences like e: one can experience a noise as coming from above a shape, as coming from underneath a shape, as coming before a shape in time, as starting when a shape appears, etc. The visual and auditory experiences which are subsumed by these multimodal experiences are always the same, but the subsuming experiences are all distinct. This shows that no single relation between unimodal experiences can account for all multimodal experiences—not logical conjunction, co-consciousness, co-introspection, simultaneity, or any other relation which has been thought to play the role of unifying experiences.

5.2.2 What is a manner of representation?

One of the main difficulties with IV- is that an interpretation of the notion of a manner of representation which makes it plausible is elusive. A manner of representation, on the most natural understanding of this expression, is simply a way of representing. Any property of a representation constitutes a way of representing, so every set of phenomenal states arguably has a characteristic manner of represen-

tation. For example, one could say that experiences in set $\{s_0,...,s_n\}$ represent in the $s_0...s_n$ way, in that they are experiences with the property of being in $\{s_0,...,s_n\}$. More importantly, there is a manner of representation all phenomenal states have: the phenomenal way of representing (i.e. representing with some phenomenal character, or representing consciously). This means that IV entails pure virtualism on the default reading: if IV is true, there is a virtual relation R corresponding to the phenomenal way of representing such that all phenomenal states are states of standing in R to certain propositions. On the default reading of "manner of representation", then, IV- is clearly false, because IV entails pure virtualism.

At the other extreme, one could take manners of representation to be the maximally specific ways phenomenal states represent. Then every manner of representation would be had by precisely one phenomenal state. On this reading, IV is equivalent to weak virtualism. This makes it highly plausible, but this makes IV- implausible in light of SV. As I mentioned earlier, it is plausible that SV generalizes. A proper generalization of SV would be at least as strong as SV, but weak virtualism is not. If there is no other reading of IV than that which makes it equivalent to weak virtualism, then the only candidate generalization of SV is the pure virtualist view.

The bottom line is that IV- is unsustainable unless it is possible to specify a technical notion of manner of representation which makes IV stronger than sensory virtualism but weaker than pure virtualism.

Sensory modalities

It is natural to begin by looking for some generalization of the concept of a sensory modality. There seems to be four reasonably natural ways of individuating sensory modalities. These four possible individuation criteria are distinguished by Grice (1962; 1988):

- Normal external stimuli, e.g. visual experiences are experiences normally caused by light.

- Characteristic representational contents, e.g. visual experiences represent colors and shapes.

- Phenomenal characteristics, e.g. visual experiences have a certain visual phenomenology.

- Organs and internal mechanisms, e.g. visual experiences are produced by a mechanism which spans the eyes, the optic nerves and the visual cortex.

We could specify these criteria more precisely, but this will do to start. For all I know, each of these ways of individuating modalities could be that which is recommended by the conventional meaning of "sensory modality" or "sense". I take no stand on this matter, and nothing in the rest of this chapter turns on this.[4] I will write $\text{modality}_{\text{sti}}$, $\text{modality}_{\text{rep}}$, $\text{modality}_{\text{phe}}$, and $\text{modality}_{\text{int}}$ for these different kinds of modality whenever disambiguation is necessary. I will also subscript modalities modifiers whenever appropriate. For example, one might say using this notation that the property of being a $\text{visual}_{\text{sti}}$ experience differs from the property of being a $\text{visual}_{\text{phe}}$ experience because the former requires a certain kind of cause but no specific phenomenology while the latter requires a certain kind of phenomenology but no specific kind of cause.

What makes non-sensory experiences non-sensory appears to be precisely that they are not caused by characteristic stimuli or produced by readily identifiable sensory mechanisms. As a result, it is difficult to make sense of the notion of a non-sensory $\text{modality}_{\text{sti}}$ or non-sensory $\text{modality}_{\text{int}}$. If manners of representation are $\text{modalities}_{\text{sti}}$ or $\text{modalities}_{\text{int}}$, non-sensory experiences don't have manners of

[4]Some claims I made earlier are not neutral regarding the nature of modalities. As noted in section 5.1.1, pure virtualism requires that we understand the modalities referred to in SV along representational lines. My defense of SV itself is not entirely neutral. For example, it does not seem plausible that any phenomenal state essentially involves any sensory organ, so it does not seem plausible that any phenomenal state is identical to $\text{seeing}_{\text{int}}$ a certain content (an implication of PCSE on the assumption that the relevant modalities are individuated by internal mechanisms). Still, I don't want to debate what a "sense" is in everyday life because this anthropological question is ultimately irrelevant to my goals. What matters is that there is an interpretation of SV along the lines of one of the preceding interpretations of "sensory modality" which is true, and that this thesis militates for pure virtualism.

representation and IV is not a generalization of SV to non-sensory experiences (it does not apply to non-sensory experiences). This would not make IV implausible, but this would make it irrelevant to our discussion, because we want to know which generalization of SV to non-sensory experiences is the correct one.

It remains to see if we can define manners of representation by reference to $\text{modalities}_{\text{rep}}$ or $\text{modalities}_{\text{phe}}$.

$\textbf{Modalities}_{\textbf{rep}}$

The $\text{modality}_{\text{rep}}$ understanding of manners of representations leads back to the problem which made us reject the traditional formulations of representationalism in chapter 1: how should we understand the relevant representational contents? We found no other satisfactory way of clarifying the notion of content than those provided by pure virtualism, sensory virtualism, and weak virtualism. Rather than explain representationalism in terms of the notion of content, we decided to do the reverse and explain the notion of content in terms of virtualism: the content of an experience is simply the entity (the relatum) which makes pure virtualism, sensory virtualism, or weak virtualism true for this experience (we must choose one of the three).[5] This does not mean that we must reject the notion of a $\text{modality}_{\text{rep}}$, but we can only understand $\text{modalities}_{\text{rep}}$ as individuated by the propositions which make specific virtualist theories true.

This stricture on content talk greatly complicates a full assessment of the $\text{modality}_{\text{rep}}$ account of manners of representation, but one plausible assumption can allow us to avoid the complications. The assumption is that, for any given experience, the identities which satisfy all the true virtualist theories for that experience are all the same. So if *e* satisfies both sensory virtualism and weak virtualism, it is in both cases in virtue of its being a state of standing in a given relation R to a given

[5] We could also say that the content of an experience is the relatum which makes IV true for this experience, but then IV would be circular if manners of representation are supposed to be $\text{modalities}_{\text{rep}}$.

proposition P. This assumption seems prima facie plausible. It is also an instance of a more general principle to the effect that identical properties have identical parts. In the case of simple relational properties or states of the form $\lambda x(R(x,\alpha))$, this means that $\lambda x(R(x,\alpha)) = \lambda x(V(x,\beta))$ entails that $V = R$ and $\alpha = \beta$. This principle implies that if a phenomenal state is both a state of standing in R to P and a state of standing R' to P', then R = R' and P = P'. In other words, the same virtual relation and proposition satisfy every virtualist theory for every phenomenal state, so every phenomenal state has at most one content.

If it is true that every phenomenal state has at most one content, the $\text{modality}_{\text{rep}}$ account of manners of representation, combined with IV, does not allow that two distinct phenomenal states have the same content. The reason is that two phenomenal states which have the same content must have the same $\text{modality}_{\text{rep}}$. On IV with $\text{modalities}_{\text{rep}}$ as manners of representation, two states which have the same $\text{modality}_{\text{rep}}$ are constituted by the same virtual relation. So, on IV with $\text{modalities}_{\text{rep}}$ as manners of representation, two states which have the same content are both states of standing in the same virtual relation to the same proposition (so are identical). IV therefore rules out distinct phenomenal states with the same content on the $\text{modality}_{\text{rep}}$ account of manners of representation. As I noted earlier, IV must allow for this possibility in order for IV- to be plausible; otherwise, there would be no perceptible advantage to IV over the simpler pure virtualist theory. IV- is therefore unacceptable on the $\text{modality}_{\text{rep}}$ account of manners of representation.

One might question the assumption that every phenomenal state has at most one content. That is, one might think that a phenomenal state could be both a state of standing in R to P and a state of standing in R' to P', where P is distinct from P'. Analogous problems to that just described still arise without the same-content assumption, but they are harder to bring out. A little bit of nomenclature is required to discuss the problems efficiently.

For any given partition Γ of a set of phenomenal states, let Γ-virtualism be

defined as follows:

Γ-virtualism For any *s* in Γ, there is a relation R virtual with respect to basic phenomenal states such that for any *e* in *s*, there is a proposition P such that *e* = standing in R to P.

Weak virtualism, sensory virtualism, and pure virtualism are all instances of Γ-virtualism. They are instances of Γ-virtualism in which Γ is the set of all singletons of phenomenal states, the set which sorts sensory phenomenal states by their sensory modalities (on some understanding of modalities), and the set which includes only the set of all phenomenal states, respectively.

We can define the Γ-*content* of a phenomenal state *e* as the proposition which makes Γ-virtualism true for *e.* That is to say that *e* has Γ-content P just in case P is the unique proposition such that Γ-virtualism is true and

1. there is a set S in Γ such that *e* is in S and

2. there is a relation R virtual with respect to basic phenomenal states such that for any phenomenal state *e*' in S there is a proposition Q such that *e*' = standing in R to Q and

3. *e* = standing in R to P.

If we want to allow more than one content per phenomenal state, we need to understand $\text{modalities}_{\text{rep}}$ in terms of specific content types: we have to speak of Γ-$\text{modalities}_{\text{rep}}$ rather than $\text{modalities}_{\text{rep}}$ tout court. For example, we could say that manners of representation are Γ_{WV}-$\text{modalities}_{\text{rep}}$, where Γ_{WV} is the set of singletons of phenomenal states (the set Γ such that weak virtualism = Γ-virtualism). Γ_{WV}-$\text{modalities}_{\text{rep}}$ are modalities which are determined by the propositions which make weak virtualism true for specific experiences. We must now ask whether there is a set Γ such that Γ-virtualism is plausible and the manners of representations referred to in IV and (indirectly) in IV- can be understood as Γ-$\text{modalities}_{\text{rep}}$.

An obvious problem arises in the case where Γ contains the set of all phenomenal states. Since Γ partitions the set of all phenomenal states, it can only contain one set in this case: the set of all phenomenal states. So, if Γ contains the set of all phenomenal states, it is Γ_{PV} and we individuate modalities by Γ_{PV}-contents (the contents which satisfy pure virtualism). If IV's modalities are individuated in this manner and it is non-vacuously satisfied, pure virtualism must be true, because the mere fact that some phenomenal state has a manner of representation (a Γ_{PV}-$\text{modality}_{\text{rep}}$) entails that all phenomenal states satisfy pure virtualism. This makes IV- untenable. There is an equally obvious difficulty if we take Γ to be Γ_{SV}: since sensory virtualism only applies to sensory phenomenal states, non-phenomenal states do not have Γ_{SV}-$\text{modalities}_{\text{rep}}$.[6]

The problem with IV- if we take manners of representation to be Γ_{WV}-$\text{modalities}_{\text{rep}}$ is that it does not allow experiences with the same content (i.e. the same Γ_{IV}-content) to differ in phenomenal character, contrary to the requirement explained on page 146. The reason is that if an experience has P as Γ_{IV}-content (on the Γ_{WV}-$\text{modalities}_{\text{rep}}$ account of IV's manners of representation), it also has it as Γ_{WV}-content. Since an experience's Γ_{WV}-content determines its Γ_{WV}-$\text{modality}_{\text{rep}}$, it is not possible for two distinct experiences to have the same Γ_{IV}-content: if they have the same Γ_{IV}-content, they also have the same manner of representation and involve the same virtual relation, so they are identical.

The key claim here is that if an experience has P as Γ_{IV}-content, it also has it as Γ_{WV}-content. If *e* has any Γ_{IV}-content at all and the manners of representation referred to in IV are Γ_{WV}-$\text{modalities}_{\text{rep}}$, *e* satisfies all the conditions listed above for having a Γ_{WV}-content (otherwise it wouldn't have a manner of representation).

[6]Of course, we could say that the manner of representation of a phenomenal state is its Γ_{SV}-$\text{modality}_{\text{rep}}$ if any and some catch-all manner M* otherwise. Then every phenomenal state would be guaranteed to have a manner of representation. But this seems rather ad hoc. Besides, there is as much a need to distinguish between the various manners of representations of non-sensory phenomenal states as there is a need to distinguish between the manners of representation of sensory phenomenal states. In particular, if different sensory experiences have different manners of representation, then it is plausible that emotional experiences and sensory imagination also differ with respect to their manners of representation.

The question is whether the same proposition satisfies both the conditions for Γ_{WV}-content and the conditions for Γ_{IV}-content for e. The key here is that, whatever elements of Γ_{WV} and Γ_{IV} e is in, we know that the relation, proposition and set of experience which satisfy conditions 1-3 for Γ_{IV}-content for e also satisfy them for Γ_{WV}-content, because all elements of Γ_{WV} (singletons of phenomenal states) have supersets in Γ_{IV}. So if proposition P, relation R and, set S satisfy conditions 1-3 above for e and Γ_{IV}, they satisfy the same conditions for e and Γ_{WV}. Only one proposition can be the Γ_{WV}-content of e, so the Γ_{WV}-content of e is P (the same as its Γ_{IV}-content).

Note that this argument applies to all content type pairs $< \Gamma_1, \Gamma_2 >$ where all elements of Γ_1 have supersets in Γ_2. When this is the case, the same virtual relation and proposition which meet the conditions for Γ_2-content for an experience meet the conditions for Γ_1-content for the same experience, so an experience which has P as Γ_2-content automatically has it as Γ_1-content as well. This suggests an extension of the argument to Γ_{SV}-modalities$_{\text{rep}}$, because it seems that a plausible alternative to pure virtualism should be a theory on which the manners of representation of sensory phenomenal states correspond to sensory modalities or more comprehensive categories (a plausible alternative to pure virtualism should be at least as strong as sensory virtualism). If the modalities$_{\text{rep}}$ referred to in IV were defined in such a way that IV meets this condition, a sensory phenomenal state which had P as Γ_{IV}-content would also have it as Γ_{SV}-content, so it would not be possible for two distinct phenomenal states to have the same Γ_{IV}-content (because they would also have the same manner of representation).

Modalities$_{\text{phe}}$

If manners of representation cannot be defined as modalities$_{\text{int}}$, modalities$_{\text{sti}}$, or modalities$_{\text{rep}}$ consistently with IV-, it is natural to try to individuate them purely in terms of phenomenology, which would make them modalities$_{\text{phe}}$. What, then, of

the proposal that manners of representation correspond to phenomenal types?

The key question here is: what kinds of kind of phenomenal character individuate modalities$_{phe}$? The answer to this question is crucial to the viability of IV-. Note in particular that IV would imply pure virtualism if all phenomenal types characterized modalities$_{phe}$. On this understanding of IV, there would be a modality corresponding to the most general kind of phenomenal character (the kind which encompasses all phenomenal characters), and all phenomenal states would be identical to states of standing in the virtual relation associated with this kind to certain propositions, so pure virtualism would be satisfied. For IV- to be viable, there has to be a suitable characterization of the phenomenal kinds which individuate modalities$_{phe}$.

One reasonable account of modalities$_{phe}$ is that visual$_{phe}$ experiences have visual content, auditory$_{phe}$ experiences have auditory content, and so on. This is how one must think of phenomenally individuated modalities within the pure view. But this does not work within IV- if manners of representations are modalities, because IV- must allow that manners of representation float free of contents. The problem is to provide an account of modalities$_{phe}$ which does not make them a function of representational content or render IV- dubious in some other way.

The sensory modalities$_{int}$ or modalities$_{sti}$ could perhaps serve as reference points. One could say for instance that a modality$_{phe}$ is a phenomenal type of the kind that captures similarities in phenomenology of the kind we find between the phenomenal states normally caused by visual stimuli or activities in visual sensory organs. One problem with this is that there are well-known cross-modal effects in which experiences we would normally classify as belonging to one modality$_{phe}$ are caused by stimuli associated with another modality$_{phe}$. For example, Sham et al. (2000) report that a single flash presented concurrently with multiple short beeps is normally perceived as multiple flashes. According to Sham et al., the illusion is so compelling that even subjects to whom it has been explained report seeing multiple flashes. Another example is provided by the classic experiment of McGurk & Mac-

Donald (1976), in which the pairing of an auditory /ba/ with a visual experience of a mouth making the sound /ga/ results in an auditory experience of /da/. Such cross-modal effects are pervasive and systematic, particularly between the vision-audition and taste-smell modality$_{sti}$ pairs (Thesen et al. 2004; Auvray and Spence 2008; Bult et al. 2007). In addition to making it hard to assign specific modalities$_{sti}$ or modalities$_{int}$ to specific experiences, these phenomena show that experiences' modalities$_{sti}$ and modalities$_{int}$ do not systematically correspond to what we would otherwise think of as their modalities$_{phe}$, insofar as we have a grasp on the latter. As a result, it is doubtful that modalities$_{sti}$ or modalities$_{int}$ could be used to elucidate the nature of modalities$_{phe}$.

Another difficulty is that it is not obvious how to generalize from whatever examples of sensory modalities$_{phe}$ we might have. Let us say for the purposes of illustration that the visual$_{phe}$ phenomenal states are those which have visual content (they are the visual$_{rep}$ phenomenal states; most types of content would do here). We could then define modalities$_{phe}$ as sets of states which are similar to each other in the way that all visual$_{phe}$ phenomenal states are similar to each other. But visual$_{phe}$ phenomenal states thus understood are similar to each other in many respects, including their contents. Personally, I cannot see what the relevant similarity is, aside perhaps from the fact that all visual$_{phe}$ phenomenal states are states of perceiving$_{p}$ colored things, which is a fact about their contents. To clearly define modalities$_{phe}$ in this kind of way, one has to give some indication of what shared properties of visual$_{rep}$ experiences make them visual$_{phe}$ experiences other than their having contents of a certain kind.

If you think you know how to individuate modalities$_{phe}$ in a way which does not reduce them to modalities$_{rep}$, I suggest that you test your concept of a modality$_{phe}$ by asking yourself what modalities$_{phe}$ there are exactly. How many sensory modalities$_{phe}$ are there? Is there only sight, hearing, touch, smell, and taste as children's books would tend to suggest? Let us consider a few examples.

Consider first the case of pain. Are pain experiences in the same modality$_{\text{phe}}$ as touch experiences? A skin irritation feels quite similar to the touch of an abrasive surface. On the other hand, headaches generally don't feel like varieties of touch. Should we distinguish different kinds of pain, some which are touch experiences and others which are not?

Or consider the perception of hot and cold. Berkeley famously argues that an intense heat is a pain,[7] and he is certainly right that it is hard to see why we shouldn't count an experience of intense heat as a pain experience. But a mild heat experience is clearly not a pain. So it is not clear whether experiences of heat are experiences of pain. It is correspondingly unclear whether experiences of heat are in the same modality$_{\text{phe}}$ as experiences of pain. The same goes for touch experiences and experiences of heat. There is something tactile about heat perception, but is it enough to make heat experiences touch experiences?

Are smell and taste really different modalities$_{\text{phe}}$? They certainly seem to have a lot in common phenomenally when one smells a wine, but they are generally considered distinct. How are we to decide the matter?

Cognitive modalities$_{\text{phe}}$ are just as problematic as sensory modalities$_{\text{phe}}$. First, is there such a thing as a specifically cognitive modality$_{\text{phe}}$? When you look at a scene and experience a distinctive phenomenology as you apprehend that something is missing, is the phenomenology part of your sensory$_{\text{phe}}$ experience or part of a cognitive$_{\text{phe}}$ experience? When you use your sensory imagination, say when you visualize something, are you representing visually$_{\text{phe}}$ or cognitively-visually$_{\text{phe}}$, or just cognitively$_{\text{phe}}$? There is certainly a difference in phenomenal character compared to ordinary perceptual experiences, but is it a difference that implies a difference in phenomenal modality?

How many non-sensory modalities$_{\text{phe}}$ are there? Does sensory imagination involve a different phenomenal modality than other kinds of cognitive experience, say

[7] In the first dialogue of *Three Dialogues Between Hylas and Philonous in Opposition to Sceptics and Atheists.*

pure phenomenal thoughts? Do emotional feelings fall under the same phenomenal modality as bodily experiences or sensory imagination?

How are we to go about answering all these questions? Assuming that modalities$_{phe}$ are not equivalent to modalities$_{rep}$, how are we to determine which differences in phenomenal character correspond to differences in modality$_{phe}$ and which do not? This does not seem to me to be an empirical question. This is a criterial question whose answer ought to be contained in the very idea of a modality$_{phe}$. I have no clue what the answer to the question might be unless modalities$_{phe}$ are identified with modalities$_{rep}$. This indicates to me that I have no idea what a modality$_{phe}$ is supposed to be if it is not the same as a modality$_{rep}$. We will have to wait and see if someone can propose a suitable individuation criterion for modalities$_{phe}$, but the odds look slim.

Other options

I discussed four approaches to manners of representation inspired by the four standard accounts of sensory modalities. I am not aware of any other proposal. In fact, proponents of manners of representation such as Chalmers (2004) and Crane (2003) appear content with the default understanding of them ("ways of representing"). This is understandable because the default understanding works just fine within Chalmers' and Crane's theories. However, it does not work as an account of the manners of representation referred to in IV and IV-, because it results in an interpretation of IV on which IV entails pure virtualism. Since the prospects for a suitable account of manners of representation seem dim, I tentatively conclude that pure virtualism is a significantly more promising theory than IV-.

5.3 The case against pure virtualism

Several putative counterexamples to pure virtualism have been suggested which might seem to favor IV-. All such cases involve two experiences in different modalities which allegedly have the same representational content while differing in phenomenal character. I believe that these cases have received their fair share of attention in the literature, so I will keep my survey to a minimum.

5.3.1 Block's examples

Block (1996) suggests a few mundane examples of experiences in different sensory modalities which could allegedly have the same representational content while differing in phenomenal character. His best case is arguably that which involves two experiences of movement. One is an experience in which an individual hears something falling from above. The other is an experience in which an individual sees something falling from above. Block claims that two such experiences could have the same content while differing in phenomenal character (one has an auditory character, the other a visual character). Block appears to have intended this case as a counterexample to all representationalist views, but it is clear that it only threatens cross-modal views such as pure virtualism.

A common response to this case is that the auditory and visual experiences in question represent more than that something is falling from above: the auditory experience also represents auditory properties, and the visual experience represents colors and shapes. These differences, the response goes, account for the difference in phenomenal character between the two experiences.

Block (1996) responds that, while the full auditory and visual experiences have distinct contents, it is possible to isolate in imagination mere visual and auditory experiences of "something falling", and these can be seen to differ in phenomenal character.

This rejoinder has the drawback of reducing Block's case to a mere thought experiment. There is also good reason to question the intelligibility of the scenario. We are asked to imagine visual$_{phe}$ and auditory$_{phe}$ experiences which merely represent that something is falling from above. Personally, I don't know what colorless, soundless auditory$_{phe}$ and visual$_{phe}$ experiences are like; I am unable to imagine such experiences. Whenever I try to imagine a visual experience, for example, I end up imagining an experience of a shape with a particular color. That is the only positive image I can form of a visual experience. This does not show that colorless visual experiences are impossible, but this means that I cannot confirm a priori that they are possible either.[8]

5.3.2 Synaesthesia

According to Galton (1980), synaesthesia is a condition in which experiences in multiple modalities occur in response to stimulation of one modality, e.g. when visual experiences occur in response to auditory stimulation. Rosenberg (2004) claims that synaesthesia is a counterexample to representationalism. More specifically, he claims that individuals who have visual experiences as a result of auditory or pain stimuli are counterexamples to representationalism. The idea appears to be that visual experiences triggered by sounds or pain do not represent their normal objects (colored expanses) but sounds or bodily damage, respectively.

SV might seem untouched by this argument since it allows visual experiences with sound or damage content. However, SV cannot accommodate Rosenberg's apparent assumption that experiences represent the events or properties that cause

[8]I am implicitly relying on Chalmers' (2002) distinction between positive and negative conceivability. A claim is negatively conceivable iff it cannot be ruled out on rational reflection. By contrast, positive conceivability requires that one be in a position to form a positive conception of a situation satisfying the claim. Positive conceivability is strongly linked with the ability to experience a situation in imagination—it is unclear what else could constitute a positive conception of a situation except an episode of imagination or a perceptual experience. Block's scenario seems to be at best negatively conceivable. But negative conceivably is a much less reliable indicator of possibility than positive conceivability (ibid.).

them. The reason is that it is clearly possible for two individuals to have different visual experiences in response to the same stimulus. SV does not allow different phenomenal states to have both the same content and the same modality. Given that we have good reasons to endorse SV, we should reject Rosenberg's assumption that experiences represent what cause them.

It is also worth noting that this is a claim which friends of causal-informational theories of content for phenomenal states reject. On Dretske's (1995) view, for example, an experience represents what it has the biological function of carrying information about. Synaesthetic visual experiences have the biological function of carrying information about colors and shapes like normal visual experiences, so they do not differ in content. Independently of commitments to a simplistic causal-informational theory of phenomenal content, I see no reason to think that synaesthetic $\text{visual}_{\text{phe}}$ experiences do not have colors and shapes as objects. Indeed, they are naturally described as cases of seeing_{p} colors and shapes. Synaesthetic visual experiences caused by sound stimuli seem to be $\text{visual}_{\text{rep}}$ and $\text{visual}_{\text{phe}}$ even though they are $\text{auditory}_{\text{sti}}$—they seem to be perfectly ordinary $\text{visual}_{\text{phe}}$ experiences with unusual causes.

5.3.3 Facial vision

Some blind individuals have the ability to orient themselves and detect medium-sized objects in their environment using echolocation. This phenomenon was initially dubbed "facial vision" because the subjects were thought to perceive their environments through the effects of light on their faces, but better studies have shown that echolocation is used, and that the ability is spread to the entire population, though far less developed in normally sighted individuals (see Ono, Fay & Tarbell 1986 for a review).

Lopes (2000) claims that proficient human echolocators are counterexamples to representationalism. According to him, we should expect echolocation expe-

riences of shape, distance, and other properties represented in vision to differ in phenomenology from visual experiences of the same properties:

> [...] surely what it is like to hear a round, velvety object three metres away is not what it is like to (dimly) see a round, velvet object three metres away. Nor, for that matter, is the phenomenal character of an experience of hearing a triangular shape the same as that of touching a triangular shape. I conclude that Dretske's representationalism is false. (Lopes 2000)

Lopes does not provide any evidence for his claim that what it's like to "hear" an object is different from what it's like to "see" it. Here it is particularly important to keep track of the different ways of individuating modalities outlined in section 5.1.2. Since $\text{modalities}_{\text{phe}}$ are defined as modalities individuated by phenomenal character, it is trivially true that $\text{hearing}_{\text{phe}}$ a round object and $\text{seeing}_{\text{phe}}$ a round object are phenomenally distinct states (if there are such states; assuming phenomenal modalities do not overlap). However, Lopes did nothing to establish that human echolocators hear_{phe} visual properties, as opposed to merely $\text{hearing}_{\text{sti}}$ them.

There is in fact some evidence that human echolocation generates $\text{visual}_{\text{phe}}$ experiences even though it is a kind of $\text{hearing}_{\text{sti}}$. Ono, Fay & Tarbell (1986) report two relevant findings. First, the brain regions associated with human echolocation are largely the same as those normally associated with vision (and not audition). Second, The functional characteristics of human echolocation are similar to those of human vision: the perceived area of space appears to be roughly the same, and human echolocation appears to have the same kind of foveal-peripheral structure as vision. Lopes dismisses the Ono et al study as "inconclusive" and "very speculative" with regard to the phenomenology of human echolocation without addressing the specific findings reported. Ono et al do not prove that echolocation generates $\text{visual}_{\text{phe}}$ experience, but I believe pace Lopes that the findings they report provide significant support for this conclusion.

Note that this hypothesis is not in tension with the fact that blind human echolocators do not claim to "see": they know very well that they cannot see_{sti} or see_{int}, and they are not in a better position than us to tell whether the phenomenal characters or contents of their experiences are those of $\text{visual}_{\text{phe}}$ or $\text{visual}_{\text{rep}}$ experiences. The reason they do not claim to see is plausibly that they face an acute version of the problem of other minds. We assume that normally functioning individuals (vision-wise) all have visual experiences of the same kind because we assume that phenomenal types go with functional and physiological types. Absent evidence of the kind I just discussed regarding functional and physiological similarities between human echolocation and $\text{visual}_{\text{phe}}$ experience, neither the echolocators nor we would have any evidence to support the hypothesis that the former's experiences are $\text{visual}_{\text{phe}}$.

5.4 Summary

In this chapter I have argued for the pure virtualist theory based on sensory virtualism, the view defended in chapter 4. I have suggested that sensory virtualism must be generalized to all experiences. Pure virtualism is one generalization of sensory virtualism. I have suggested that it is more phenomenologically accurate than the alternative (IV-), and that the alternative faces serious difficulties. I have raised two problems for IV-: it cannot account for multimodal experiences, and it seems false unless the notion of a manner of representation it appeals to is given a technical definition which is elusive. I have concluded with a brief discussion of three putative counterexamples to the kind of cross-modal representationalism which pure virtualism exemplifies. One posits experiences which cannot be imagined, one trades on a bad theory of phenomenal content, and the other seems to rest on an equivocation between the different kinds of thing one could designate as "modalities", in addition to running afoul of the scientific evidence there is on the question.

Part III

Alternatives

Chapter 6

On the case for qualia

Virtualism analyzes basic phenomenal states into two components without remainder: relation R, and the propositional relata of R (the contents of phenomenal states). An immediate implication of this *pure propositionalism* is that the phenomenal characters of experiences supervene on their contents: any two experiences with the same content have the same phenomenal character. A number of putative counterexamples to this supervenience thesis have been suggested aside for the cross-modal cases discussed at the end of the preceding chapter. For example, it has been claimed that visual experiences of the kind we undergo when our vision is blurred can have the same contents as visual experiences resulting from clear vision even though the two kinds of experience differ in phenomenal character (Boghossian and Velleman 1989; Smith 2008). Other kinds of perceptual variation (e.g. perspective and double vision) have been suggested as sources of counterexamples to the supervenience thesis (Boghossian and Velleman 1989; Peacocke 1983). Gestalt phenomena (Peacocke 1983, Nickel 2006) and the possibility of spectrum inversion (Block 1996; 1998; 2003, Shoemaker 1994; 2000; 2001) have also been singled out. In addition to purportedly refuting propositionalism and representationalism, these cases constitute the main evidence for the existence of qualia. I will address these objections after preliminary remarks regarding the notion of content.

6.1 The relevant notion of content

As typically formulated, the preceding challenges presuppose some understanding of the notion of phenomenal content which is compatible with rejecting virtualism and all but the weakest forms of representationalism: without such a notion of phenomenal content, one cannot say that certain variations in phenomenology are independent of variations in content and sensory modality. I discussed a number of theory-independent notions of content in chapter 1 (e.g. how things seem in an experience, what information an experience carries, etc.), but none of these kinds of content are relevant to our project. How, then, are we to make sense of the present objections?

The objections apply to virtualism, but they need to be reformulated. In chapter 5, I claimed that the relation which satisfies virtualism is the experiencing relation. The experiencing relation is the salient relation that is common to seeing$_p$, hearing$_p$, and so on (p. 142). The cases mentioned above would refute the kind of virtualism I recommend (i.e. virtualism+) if they showed either a) that some phenomenal states do not involve standing in the experiencing relation to propositions or b) that there are possible distinct phenomenal states which involve standing in the experiencing relation to the same proposition.

Option (a) can be set aside at this stage. The evidence I gave for the availability of phenomenal readings of perceptual ascriptions and a propositional analysis of such ascriptions is not put in question by the cases at hand, so these cases do not throw doubt on the existence of the experiencing relation. It also seems highly plausible that the phenomenal states they refer to consist at least in part in standing in the experiencing relation to certain entities. For example, it makes no doubt that blurry visual experiences are seeings$_p$. At best, the case of blur shows that they are blurry seeings$_p$, or seeings$_p$ accompanied by certain qualia. More generally, it seems plausible that all visual phenomenal states (however anomalous) involves seeing$_p$ something. Since all the examples mentioned are visual phenomenal states,

none of them seems to challenge the claim that all phenomenal states involve standing in the experiencing relation to something. For the purposes of interpreting the present objections, we can take the content of an experience to be the entity one stands in the experiencing relation to in this experience.

The question, then, is whether two phenomenal states which consist at least in part in standing in the experiencing relation to a given content can differ in phenomenal character due to the influence of blur, perspective, spectral inversion, etc on their phenomenology. If there were such states, virtualism+ would be false and the perceptual conception of experience discussed in 3.2 would be seriously threatened. It would not be ruled out that virtualism is satisfied by some other relation than the experiencing relation, but I would see little reason to expect this.

6.2 Blurry and double vision

Boghossian and Velleman (1989) object to representationalism as follows[1]:

> [...] you can see nearby objects double by focusing on distant objects behind them, and yet you cannot get yourself to see the number of nearby objects as doubling. And by unfocusing your eyes, you can see objects blurrily without being able to see them as being blurry. None of these experiences can be adequately described solely in terms of their intentional content.

Before we address the objection, it is important to note that no experience or phenomenal state is intrinsically or essentially blurry. Blurry experiences are like blurry photographs. What makes a photograph blurry is not merely a matter of how its pixels are arranged, but also of how they are arranged compared to how things are: a picture of a fog or any other scene apt to cause a pixel pattern identical to that of a

[1] A.D. Smith (2008) also discusses blur. However, his argument explicitly rests on the assumption that the contents of experiences are accuracy conditions, which we rejected in chapter 1.

blurry photograph need not be blurry. Similarly, what makes an experience blurry is not merely a matter of what its phenomenal character is like, but also of how its phenomenal character relates to one's environment. To a first approximation, what one normally means when one says that one's vision is blurry is that one's visual experiences are distorted in a certain way compared to the norm. In some cases, this is because one's visual apparatus is malfunctioning. In other cases, one's visual apparatus can be functioning correctly while receiving input from another apparatus (e.g., lenses, a camera, or a film) which distorts the final experience in a relevant way. Since blurriness is a contingent characteristic of visual experiences,[2] it is misleading to speak of "blurry experiences" in the same way that we speak of shape and pain experiences, which are essentially related to shapes and pain, respectively. To avoid confusion, I will refer to experiences occurring in blurry vision as *b-events.* I will refer to the phenomenal states instantiated in blurry vision as *b-states*.

Parallel remarks apply to double vision. To suffer from double vision is not to have experiences with a certain kind of phenomenal character, but to have experiences whose phenomenal characters are distorted in a certain way compared to those of experiences one would normally have when looking at the same scene. We can refer to experiential events which are part of episodes of double vision as *d-events* and phenomenal states which are part of episodes of double vision as *d-states*. I will refer to experiences and phenomenal states which do not fall in the b- or d- category as *normal experiences* and *normal phenomenal states*, respectively.[3]

Boghossian and Velleman appear to be suggesting the following argument from blurry vision:

Boghossian and Velleman's argument (blurry vision)

1. One does not always see anything as blurry when one is in a b-state.

[2]Here it is important to keep in mind that we are individuating experiences by their phenomenology.

[3]Note that the preceding definitions of b-, d- and normal phenomenal states are non-rigid. These qualifiers merely characterize how the state is produced in a given world.

2. If (1), b-states cannot be described completely in terms of their contents.

3. If b-states cannot be described completely in terms of their contents, representationalism is false.

   Therefore, representationalism is false.

Premise 1 has two natural readings:

1. One does not always believe that anything is blurry when one is in a b-state.

2. One does not always see_p anything as blurry when one is in a b-state.

Subtle variations on these readings are possible, but these appear to cover the relevant possibilities.

On the first reading, premise 1 seems undeniable for the simple reason that "blurry" is an adjective which does not apply to most of the objects we see when we see blurry. Normally, we don't believe that the things we are looking at when our vision is blurred are blurry because that would be a category mistake—only the products of optical processes can be blurry. While premise 1 is undeniable on the first reading, premise 2 is questionable. In chapter 1, we went to great lengths to divorce phenomenal contents from propositional attitudes. I agree that the phenomenal character of experience is generally not reflected in full in the beliefs one has about the world, but I take this to show merely that phenomenal content floats free from belief content, not that phenomenal character floats free from phenomenal content. There is nothing in our notion of visual phenomenal content—what one $sees_p$—to suggest that visual experiences' contents are limited by the contents of beliefs.

The second reading of premise 1 also makes it undeniable. We should not expect the term "blurry" to always apply to the kinds of things we see_p when we see blurry, because we do not generally see_p these things as products of optical processes, and only a product of an optical process can be blurry. However, Boghossian and

Velleman's platitudinous observation in premise 1 does not on the face of it bear on the consequent of premise 2. Why could not the contents characteristic of b-states be contents which involve entities having properties other than blurriness? It is certainly not true in general that if an experience can be described as F and the experience's phenomenal character is exhausted by its content, then property F is part of the content of the experience. For example, one of my visual experiences this morning resembled the latest Jim Carrey movie in some ways (anything resembles anything in some respects, after all). Clearly, propositionalism does not require that the property of resembling the latest Jim Carrey movie in some ways was part of the content of my experience. In general, contingent properties of experiences need have no echo in their contents.

The argument from double vision goes as follows:

Boghossian and Velleman's argument (double vision)

1. One does not always see anything as doubled when one is in a d-state.

2. If (1), d-states cannot be described completely in terms of their contents.

3. If d-states cannot be described completely in terms of their contents, representationalism is false.

   Therefore, representationalism is false.

Here too there are two reasonable readings of "see anything as". If Boghossian and Velleman mean see_p, then premise 1 seems dubious, because we naturally describe d-states as states in which we see_p multiple copies of a single object ("How many fingers do you see?"). On the other hand, if they mean something like "you cannot get yourself to believe that the number of nearby objects is doubling", then the point is irrelevant to virtualism. Of course we take double vision to be a kind of hallucination or illusion; we don't endorse the contents of these experiences. This does

not the least suggests that the consequent of premise 2 is true, because phenomenal content is not tied to belief content.

A number of theorists take blur and double vision to be major challenges for representationalism even though they would not necessarily endorse exactly Boghossian and Velleman's objections. I now want to suggest a general argument against the possibility of normal a phenomenal state having the same content as a b-state.

Consider the two series of experiences illustrated in figure 6.1. The images of

Figure 6.1: B-state series

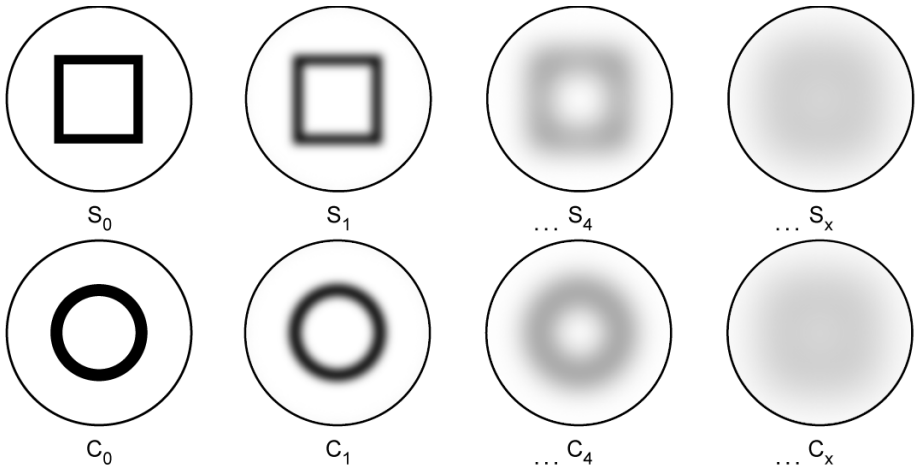

the S series are analogous to the phenomenal characters of b-states which could be produced by looking at a square while one's vision is becoming increasingly blurry. Those of the C series are analogous to experiences of a circular shape under parallel circumstances.

S_0 and S_x clearly differ in content: S_0 is a state in which one experiences and $sees_{ip}$ a square, while S_x is not. This is particularly clear in light of the fact that a) C_x and S_x could very well have the same phenomenal character (hence be the same phenomenal state) and b) there is no reason to think that C_x is a state in which one $sees_{ip}$ a square.

Given that S_0 and S_x differ in content, there must be some difference in content between some consecutive pairs of experiences in the S series. Where should we

locate the change? The only remotely plausible answer is that each two consecutive experiences in the series differ a little bit in content. But this would rule out that S_0 (the normal state) has the same content as any other phenomenal state in its series (as any b-state in the series).[4] We can construct similar gradual series starting from any normal experience and passing through any b-state. For each such series, we will have to postulate that the contents of the relevant phenomenal states differ at each step in order to explain how the starting point and the end point can differ in content. This strongly suggests that no normal state ever has the same content as a b-state.

A parallel argument applies to double vision. We can easily construct series of d-states analogous to the S and C series. Such series, like the S and C series, must contain consecutive experiences with different contents. All pairs of consecutive experiences are equally good candidates, so it is reasonable to conclude that all consecutive experiences differ in content, which in turns strongly suggests that double vision always makes a difference in content.

It is worth stressing that blur and double vision need not make a difference to the beliefs we form about the external world even if they make a difference to the contents of the experiences we base our beliefs on. If my vision suddenly became blurry, I would not revise my conception of the space around me for that: I would continue to believe that the objects in front of me have sharp edges at specific locations. I would simply disregard certain aspects of what my experience tells me about my environment. The content of my experience is given by what I see_p as I undergo this experience, not by what I am inclined to believe.

Note also that the preceding argument is entirely independent of any account of the precise nature of the characteristic contents of b-states. Dretske (1995) suggests that b-states represent fuzzy properties. Tye (2002b) holds that they have indeter-

[4] Strictly speaking, this remains a logical possibility. For example, the contents of S_0 an S_2 could be the same even though S_1 differs in content from both S_0 and S_2. But I take it that this possibility can be excluded without argument.

minate contents, in the sense of contents which do not specify exactly how things are (where the boundaries of objects are, for example). I suspect that both accounts are true of certain kinds of b-state, but the matter is difficult to adjudicate. I am inclined to think that, more often than not, none of the terms so far employed to describe the contents of b-states are exactly right. I am at peace with these practical difficulties: I don't see why we should expect it to be easy to put in words all the strange and subtle variations of perceptual experience. Public language is meant to capture important, publicly observable facts. The subjective effects of blurry vision, double vision and other types of perceptual distortion are not public observable and generally not important to us. It could conceivably be a requirement on a successful language and conceptual scheme that it makes us largely blind to such perceptual variations and enforces efficient thought and communication by not providing us with any means of stating the sorts of exotic states of affairs we experience when our perception of the world is distorted.

6.3 Perspective

Peacocke (1983) argues that variations in perspective do not always correspond to variations in content. One example he gives highlights the effect of distance:

> Suppose you are standing on a road which stretches from you in a straight line to the horizon. There are two trees at the roadside, one a hundred yards from you, the other two hundred. Your experience represents these objects as being of the same physical height and other dimensions; that is, taking your experience at face value you would judge that the trees are roughly the same physical size ... Yet there is also some sense in which the nearer tree occupies more of your visual field than the more distant tree. (Peacocke 1983, ch. 1)

The idea is that one's experiences of the two trees on the road, e_1 and e_2, could have the same content even if they had different phenomenal characters corresponding to the different apparent sizes of the trees. This would contradict representationalism and propositionalism. Peacocke runs similar arguments based on other kinds of variation in perspective.

Peacocke's examples clearly bring out the fact that what we are presented with in vision are not mere three-dimensional scenes in Euclidean space or two-dimensional pixel arrays. But these are not the only possibilities.

To see that there are some differences in phenomenal content corresponding to variations in perspective, we need only consider series of variations in perspective analogous to the series of b-states discussed above. Imagine a series T_0 .. T_x of experiences of a tree seen from increasingly far away. Setting aside the phenomenology associated with the surrounding landscape, one's experience of the tree would in the limit case (T_x) be phenomenally identical to an experience of a simple dot, or something close to this. Clearly, this experience would not be a state of seeing_{ip} a tree or tree shape (however, it could be a state of seeing_{m} a tree). By the virtualist definition of phenomenal content, it follows that it is not an experience with tree content. But T_0 would plausibly have tree content. At what point do experiences in the T series stop having tree content? The only plausible answer is that each change in distance affects what one sees_{p} or represents a little bit: what one represents progressively changes from something that is definitely a tree shape to something that is definitely not a tree shape, with intermediaries that may well be neither (due to the vagueness of the term "tree shape"). Since the same reasoning applies to experiences of any shape as well as other kinds of variation in perspective (e.g. variations in viewing angle), it seems reasonable to conclude that all phenomenal differences which correspond to differences in perspective correspond to differences in content.

While there clearly seems to be differences in content which correspond to differences in perspective, it is hard to tell what these differences are exactly. Lycan

(1996b) suggests that visual experiences have multiple layers of content. An experience could, for instance, represent objects in objective three-dimensional space at one level and objects in an egocentric space at another level. Lycan suggests that the two experiences in Peacocke's example (e_1 and e_2) could share a layer of content (the objective one) while differing in another layer of content. Tye (1996) suggests that the difference in content between e_1 and e_2 is one in situation-dependent properties such as *being large from here*. There are many other candidate accounts of the phenomenal character of perspective. For example, Thomas Reid held that perceived visual space has a spherical geometry (Reid 2000 ch 6, section 9; see also Van Cleve 2002). Kant and most theorists of Reid's time disagreed with him, but nowadays the consensus among psychophysicists and psychologists appears to be that visual space is even more exotic than Reid thought. Lunenberg's (1950; 1947) and Blank's (1958) theory that perceived visual space has a hyperbolic geometry made consensus for most of the second half of the 20th century (see Wagner 2006). Recently, even the widespread assumption that visual space has three dimensions and a constant curvature has been challenged (see French 1987, Wagner 2006).

Today, none of the aforementioned theories seems to dominate. There is, it seems to me, a deep difficulty with characterizing the perspectival aspect of visual experience. As in the case of blur, the problem seems to be that our concepts are geared toward describing a perspective-independent world. We are wired to systematically look past the variations in phenomenology caused by perspective. The result is that we simply cannot articulate how the world is given to us in visual experience. At most, we can say that we are presented with *that kind* of state of affairs. Again, I don't think this should deter us from propositionalism. It is clear in light of the argument given above that variations in perspective correspond to variations in content. The root of the problem is that we are unable to characterize the relevant variations in phenomenology—it is not that the theory is false, but that the explanandum is elusive.

6.4 Gestalt effects

Gestalt effects are another common source of objections to propositionalism. I will begin with Nickel's (2006) objection from grouping effects. Nickel illustrates grouping effects with figures 6.2A and 6.2B.

Figure 6.2: Grouping effects

A

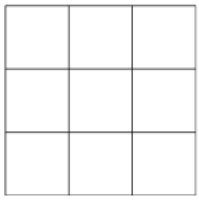

B

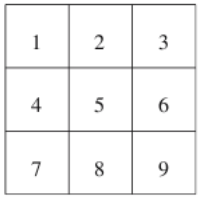

He explains the effect as follows:

> In [figure 6.2A], you can see the squares corresponding to 1, 3, 5, 7, and 9 as prominent, or you can see 2, 4, 6, and 8 as prominent. You may also be able to see other groupings as relatively more prominent, such as 1, 2, 3, 5, and 8, which form a ''T''. Consider two successive viewings of [figure 6.2A]. Suppose that you see one grouping of tiles during the first viewing, a different grouping during the second. You can have these different experiences without changing where you look. For instance, you can continue to focus your vision on the center of figure 6.2A and still have the different experiences. (284)

Nickel then argues that differences in phenomenology which correspond to differences in what group appears prominent do not correspond to differences in content between the relevant experiences. For the purposes of illustration, let us compare two specific experiences. In e_1, squares 1, 2, 3, 5, and 8 are seen as grouped in a "T"

shape. In e_2, squares 2, 4, 5, 6 and 8 are seen as grouped in a cross shape. Nickel's argument proceeds in two steps. He first argues that the sets of possible worlds at which e_1 and e_2 are true are exactly the same. He refers to this claim as *PW*. He then argues that if this first point is correct, e_1 and e_2 do not differ in content. He refers to this conditional claim as *Sameness*.

Nickel does not offer a systematic argument for PW. Instead, he argues for PW by ruling out various accounts of how the sets of worlds at which e_1 and e_2 are true might differ.

Given our notion of phenomenal content, there is an obvious account of the difference in content between e_1 and e_2. We naturally describe the difference between e_1 and e_2 as a difference in the groups one *sees* the squares as belonging to in these experiences (a T-shaped group in one case, a cross-shaped group in the other). If one sees something as an F, then one $sees_{ip}$ something being F. What one $sees_p$ is what one's experience represents (on the conception of phenomenal content we adopted). Therefore, e_1 and e_2 differ in content, because one represents the squares as belonging to different groups in these experiences.

But Nickel anticipates this reply. In response, he asks the representationalist to explain what a group is exactly. He suggests that the relevant notion of group would have to be explicated in terms of prominence. He then argues that this leads to difficulties for the representationalist because prominence must in turn be explained in terms of effects on a perceiver.

It seems to me that the groups in question might be thought of more simply as solid objects or similar entities: when we see squares as grouped, we see them as forming solid objects. That is to say that we see them as bound together in the same kind of way that a solid object's parts are bound together. So one sees a T for real—a continuous, cohesive, T-shaped lump of something. On this account, the contents of e_1 and e_2 are not true at the same worlds: the content of e_1 requires that a solid, T-shaped arrangement of squares exists but not a cross-shaped one, while that of e_2

requires that a solid, cross-shaped arrangement of squares exists but not a T-shaped one. This proposal is not considered by Nickel, and it clearly avoids the difficulties associated with the prominence account Nickel discusses.

Peacocke (1983, ch 1) also discusses the phenomenology of grouping. Instead of squares, he uses the example of rows of dots, but the considerations are essentially the same. Rather than positing grouping sensations as Peacocke suggests we must, we can explain the phenomenology of dot groupings in terms of ascriptions of ties between the dots (invisible ties holding the dots together in formation).

Peacocke also discusses the Necker cube effect. Here matters are complicated by the elusiveness of the phenomenology and content of perspective. However, there is at least one thing we can say about this effect: as the perspective from which we see the cube seems to change, we see different parts of the cube as being being located at the front. Given the relation between seeing as and phenomenal content noted earlier, we may reasonably conclude from this that the relevant gestalt shifts correspond to differences in which part of the cube is represented as being in front. It remains to explain what "in front" means in this context. This is a potentially difficult task, but one might reasonably expect a solution along the lines of Peacocke's (1992) own. That is, one might expect that relations such as *in front of*, *behind*, *to the left of* (etc.) between objects of experience can be precisely defined in terms of an egocentric space.

6.5 Cognitive experience

So far I have said little in defense of the applicability of virtualism to non-perceptual experiences (i.e. cognitive and emotional experiences) aside for the general arguments against manners of representation in the preceding chapter. It has been suggested that non-perceptual experiences are different in kind from perceptual experiences, in that the difference between these two kinds of experience lies in something

else than their contents. It is certainly true that perceptual and non-perceptual experience differ in many respects aside for their contents. For one thing, the former normally occur in the process of perceiving$_c$ external objects of certain kinds, but not the latter. But I don't think there is any difference in the phenomenal states instantiated in the two kinds of experience beyond the contents one is related to in them. This is what virtualism requires, and this is what the arguments against IV-presented in chapter 5 tend to suggest. In this section I will discuss the case of cognitive experiences in more detail. I will leave emotional experiences for the next section.

Cognitive experiences should be distinguished from other kinds of mental state which are sometimes described as "conscious thoughts". There is a common use of "conscious thought" on which "conscious" has roughly the same meaning it has in "conscious effort". In this sense, a conscious thought is roughly a thought that one has deliberately or in such a way that makes the fact that one is having it in some sense accessible to oneself. While there might be a kind of phenomenology typically associated with having conscious thoughts in this sense, conscious thoughts normally involve more than phenomenal states, because one's awareness of one's own mental state, and the mental state one is aware of being in, need not be pure phenomenal states. Conscious thoughts in the present sense are not all phenomenal states. "Conscious thought" is also sometimes used specifically to describe cognitive states which have a phenomenal component. Such states are not all phenomenal states either, because having a phenomenal state as a component is not the same as being a phenomenal state. Here I am only concerned with pure phenomenal states. By "cognitive experience" and "cognitive phenomenal state", I mean experiences and phenomenal states which are not tied to sensory or emotional processes in normal conditions. I will also refer to cognitive phenomenal states as "phenomenal thoughts". By "perceptual experience" and "perceptual phenomenal state", I mean experiences and phenomenal states which are tied to sensory processes in normal

conditions.

How much cognitive phenomenology there is is a difficult question, but there are certain kinds of cognitive experience whose existence can hardly be disputed. In particular, it seems undeniable that episodes of sensory imagination have characteristic phenomenal characters. There is a characteristic phenomenology associated with visualization, and also one associated with talking to oneself in one's head or feeling one's movements ahead of performing them. Cognitive impression*s* such as the feeling of deja vu or the feeling of confusion are equally important to the ecology of the mind. There might also be a phenomenology of thought beyond imagery and general feelings, but that is somewhat less clear. I will try to work around this uncertainty in what follows.

While the phenomenal characters of cognitive experiences differ markedly from those of perceptual experiences, it has been suggested that cognitive experiences can have the same contents as perceptual experiences. If this were the case, virtualism would be false: we would have to posit different virtual relations for cognitive and perceptual experiences. Here is an example due to Neander (1998):

> [...] as I am writing now, I remember that my neighbor has ginger hair (and I am aware of that memory). But I do not have vivid color imagery, and there is a clear phenomenal difference in the quality of the conscious experience involved in my seeing as opposed to my remembering the color of his hair.

David Chalmers has suggested to me an example which is potentially clearer. Imagine that you are looking at a simple scene, say a white wall marked with a red dot. You then close your eyes and try to think about the scene. It might seem that you could have a cognitive experience whose content encompasses the scene in as much detail and precision as your prior perceptual experience, but which has a different phenomenal character.

It is important to note that the preceding examples constitute no objection to

virtualism unless they are supposed to involve *phenomenal thoughts* with the same contents as perceptual phenomenal states. That one can entertain the same contents in thought as in perception is not an objection to virtualism, because virtualism allows that non-phenomenal thoughts or propositional attitudes are not states of standing in relation R to propositions. In order to make the case against virtualism, one must show not merely that the contents which are entertained in perceptual experience can be entertained in thought, but that they can be entertained as part of *phenomenal* thoughts. It seems far from obvious to me that this is the case in the examples suggested by Neander and Chalmers.

The limitations on recognitional abilities highlighted by perceptual recognition tasks support the view that phenomenal thoughts never represent the same contents as perceptual experiences. If one were capable of forming phenomenal thoughts about perceptible properties (properties perceptual experiences represent), one should be able to select such a property in thought prior to experiencing it perceptually, then go on to recognize it when one experiences it perceptually. This is to be expected because the primary function of thought is to determine our reactions to sensory stimuli. Of course, abstract thoughts (thoughts about social systems, for example) do not directly determine our reactions to sensory stimuli. But they determine our reactions to sensory stimuli through their causal and/or inferential connections with thoughts that have observable contents. Thoughts with observable content, including phenomenal thoughts of the kind we are considering, should connect with perceptual experiences, in that when one has a perceptual experience with the same content as a thought or a logically related content, one should be able to tell that this is the case. If this is correct, the hypothesis that we can have phenomenal thoughts with the same contents as perceptual experiences suggests that one should as a general rule be able to select the precise properties represented in perceptual experiences in thought and go on to recognize them in perception. In practice, however, we are incapable of doing this. Familiar experiences in the art class or at the

paint shop show as much. Rigorous tests which bear on this matter have also been conducted for both the recognition of color and pitch: in either modality, our ability to recognize properties is poor compared to our ability to discriminate them perceptually, and this even when the recognition task makes use of short-term memory (as opposed to long-term memory).[5] This discrepancy between the grain of perceptual representation and the grain of recognitional abilities strongly suggests that we cannot form phenomenal thoughts about perceptible properties as fine-grained as those perceptual experiences represent.

This conclusion also seems warranted by reflection on the phenomenal character of imagery. If I close my eyes and try to visualize the Coke can on my desk with as much precision as possible, the phenomenal character of the resulting experience seems to give me no hint as to the precise shade of the can. I seem to have no basis for saying that it is red_{12} rather than red_{13}, even though these are shades I can discriminate in perception. The point here is not that I cannot *name* the precise shade of red found on the Coke can. The point is that I cannot determine which precise shade of red it is based on my mental image of it. This strongly suggests that imagery represents coarser properties than those represented by perceptual experience. It does not seem implausible to say that this is all the difference in "vividness" between imagery and perceptual experience amounts to.

While the last argument bears primarily on the relation between imagery and perceptual experience, it can be extended to phenomenal thought generally. To license extending the conclusion to phenomenal thought generally, we need only remark that phenomenal thoughts of the imagistic kind appear to have the most precise contents of all phenomenal thoughts. This can be seen from the fact that someone asked to form as precise a thought as possible of a given object without looking at it will invariably engage in imagery. This suggests that phenomenal thoughts generally cannot represent the sorts of highly determinate properties represented in

[5] See Nilsson and Nelson (1981), Malone and Hannay (1978), Raffman (1995).

perceptual experience if imagistic thoughts cannot.

There seems to be good reason to deny that phenomenal thoughts relate us to the same contents as perceptual experiences, but we still have to explain the initial plausibility of the position expressed by Chalmers and Neander. One possible explanation lies in the fact that a) we can easily form *non-phenomenal* thoughts with the contents of perceptual experiences and b) for most purposes, it does not matter whether the relevant contents are grasped through experience or not, so we tend to conflate these ways of grasping content.

In section 2.4 I sketched an account of how phenomenal states realize propositional attitudes. I distinguished two kinds of attitude: conscious attitudes and external attitudes. Conscious attitudes consist in having suitably related occurrent and potential experiences. External attitudes consist in having experiences which stand in suitable relations among themselves and to non-experiences (e.g. entities in one's environment). Conscious attitudes can share content with perceptual experiences without being phenomenal states themselves, but the illusion that we can entertain phenomenal thoughts with the contents of perceptual experiences owes more to the ease with which we token external attitudes with such contents.

We have multiple ways of talking and thinking about precise colors and shapes in absence of relevant perceptual experiences, occurrent or potential. First, we have conventional means of referring to them. We have names, codes, and reference objects (*the same yellow as on the place mats*). This allows us to talk about more specific properties than we can grasp consciously in absence of relevant experiences. With this ability to talk comes an ability to think, because we think in good part through internal talk: as a general rule, when you say something to yourself, you thereby think what your words mean. When these means of reference fail, there remains the more basic method of referring to properties and objects as the causes of our experiences. For example, if I visualize an object I just observed, I am inclined to refer to the object as *that object which caused this memor*y (the residual mental

image I have of the object). I might not say this to myself, but that is nevertheless the referential intention which grounds my thought about the object, because it is the criterion I would ultimately use to identify the object: I would try to find out which object caused my experience. (Note that I am not saying that mental imagery is descriptive: my imagery experience does not have a descriptive content. It is the thought I tend to form about my mental image which has descriptive content.)

To illustrate, suppose I am looking at what seems to me to be a plant. I then leave the site of this experience and go on to talk with others about the plant. If someone were to ask which plant I was looking at, how would I go about answering their question? I would normally begin by revisiting the place in imagination. I would bring back a mental image of the place and describe the scene as I saw it: "I mean the large plant with the yellow flowers in the corner of the room, next to the portrait". If someone were to ask me what room I mean and I didn't have any conventional means of referring to it (e.g. "room 43"), I would tell them what I know of the room's relation to other parts of the world they themselves might be able to locate, and I would typically do this by navigating through a mixture of visual and motor memory: "I remember I made a right upon coming out, walked all the way down the corridor, got down to the ground level and followed the corridor on the left." So far, causation has not entered the picture yet except indirectly through my description of a procedure to get from the room to another point. That is because we generally assume that our experiences are (in relevant respects) veridical, so we give reference-fixing descriptions which are formulated directly in terms of external objects as presented by our experiences ("the plant in the corner"). But these descriptions abbreviate our true referential intentions. This comes out in cases of illusion. If I were to find out that I had not been in a proper room but in a small locker which had given me the impression of a large room through some complex trickery, I would no longer explain what I mean by reference to the room, its location, or its contents. I would say, "the plant I saw_c, if that was real." There I mean

(approximately) the plant that caused my experience of a plant-shape.

The important point is that there are different levels of thought and meaning. First, there is what we grasp directly through the experiencing relation, mostly through perceptual experience and imagery. Then there is a much larger sphere of content on which we can latch indirectly in virtue of the relations our experiences bear to other entities. Grasping things consciously is one among other means of picking things out, and what normally matters is what our words pick out. This is why we don't normally pause to distinguish between experientially grasped and non-experientially grasped contents.

6.6 Moods and emotions

Emotional states (moods and emotions) are another source of objections to representationalism. It is important to distinguish between emotional states and the phenomenal states instantiated in the latter, which I will refer to as "emotional feelings". As a general rule, moods and emotions are not phenomenal states. For one thing, they persist in absence of phenomenology. They can be unconscious in every sense of the word. They also essentially involve cognitive and behavioral dispositions. For example, one cannot be angry or in a bad mood unless one is to some extent disposed to do something mean. This does not appear to be the case of emotional feelings, the pure phenomenal states which tend to accompany moods and emotions.

Of course, here I am merely expounding features of what I take to be the folk concepts of moods and emotions. Given that there are debates about these matters, I should say that nothing of what I have to say here turns on what moods and emotions are for the folk. If the folk use these terms in such a way that they are phenomenal states or essentially involve phenomenal states, this is fine with me. However, that is not how I use the terms "mood" and "emotion". I trust that you can make sense

of how I use them sufficiently to see what I mean by "emotional feelings", which are the only states of interest here.

The main objection from emotional states against representationalism is that the phenomenal states associated with moods have no intentional content (Searle 1983; Kind 2007). Although this is the main objection, I will start with the case of emotions because it sheds light on the case of moods.

Emotions such as fear and happiness are generally agreed to have intentional content. Almost everyone agrees that they are directed at objects or states of affairs. They might nevertheless seem to pose a problem for the virtualist view, because the kinds of content they are generally agreed to have are not plausible relata of R. Take the emotion of fear, for example. Most would agree that fearing a spider is an intentional state whose content is the spider or a proposition involving the spider. The problem is that it is implausible that I stand in the experiencing relation to the big huntsman spider on the wall which scares me: there is no such thing as the phenomenology of fearing Bob the huntsman or even a spider of this specific species. There is a sense in which my fear has Bob as intentional object (on some understanding of "intentional object"), but the phenomenal state I instantiate as I am scared of Bob is not a state of standing in R to a state of affairs involving Bob or huntsmanhood. So the kind of account of the intentional contents of emotions which is generally agreed upon cannot serve as an account of the relata of emotional feelings. This does not mean that virtualism is false, because emotions are not the same as the emotional feelings which often accompany them. However, this means that we must look elsewhere for a plausible account of the contents of emotional feelings.

Such an account is to be found in a proper characterization of the phenomenology of emotional feelings. Here I think William James' well-known remarks on the nature of emotions are helpful. James appears to take emotions to essentially involve experiences. I am inclined to disagree with James on this score, but I find his

characterization of the experiences associated with emotions perspicuous. Here is how he describes some of the experiences which typically accompany the emotion of fear:

> What kind of emotion of fear would be left if the feeling neither of quickened heart beats nor of shallow breathing, neither of trembling lips nor of weakened limbs, neither of goose flesh nor of visceral stirrings, were present, it is quite impossible for me to think. (1890: 451)

While James might well be wrong about the nature of emotions, his characterization of the feelings which accompany them seems right to me: they seem to be very similar to bodily experiences. It does not seem implausible to say that they are states of standing in the experiencing relation to sensations-involving propositions, where sensations are entities of the sort we talk about when we say such things as "I feel something in my toe" (see section 3.1). The feelings of quickened heart beats, shallow breathing, trembling lips and weakened limbs James refers to are in the first instance feelings of bodily sensations similar in kind to those we feel when in pain. The same applies to the feelings associated with joy, anger, and pretty much any other emotion.

Turning now to moods, it seems plausible that the little phenomenology there is associated with moods is an attenuated version of the kind of phenomenology associated with emotions. When I am in a good mood, I feel relaxed, lighter, and generally more prone. These are primarily bodily experiences. We might not want to classify all phenomenal states associated with moods as bodily, but this does not change the fact that they are experiences of sensations of a kind with the sensations we experience in typical bodily experiences.

6.7 The inverted spectrum

Imaginary cases involving inverted spectra have received a great deal of attention in the debates surrounding representationalism. I will consider only the main case: Shoemaker's normal inverts. Block's (1996; 1998; 2003) Inverted Earth case is explicitly directed at externalist representationalism and clearly ineffectual against narrow representationalism. Since virtualism is neutral on the narrow/wide question, we need not worry about Inverted Earth. Arguably, this scenario is also easily accommodated by teleological variants on wide representationalism (see Lycan 1996c).

Shoemaker (1994; 2000) argues that two individuals or groups of individuals could be color inverted with respect to each other without either misperceiving colors, even if they perceived the same objects in the same conditions. For example, it could have been that men experience red when women experience green and men experience yellow when women experience blue. Surely, the argument goes, neither group could then be said to see colors correctly while the other does not. So, in this possible world we are describing, either both men's and women's experiences are veridical or both men's and women's experiences are not veridical. Shoemaker opts for the first option. But, assuming as Shoemaker does that different color experiences have different contents, this requires that color experiences represent relational properties such as *causing an experience of such and such type,* properties which Shoemaker calls "appearance properties" or "phenomenal properties".

Shoemaker's solution to the puzzle is not open to us, because it is prima facie implausible that we see_{p} appearance properties: we do not see_{p} experiences or experience-involving states of affairs in visual experience. Shoemaker's view is not compatible with our individuation criterion for phenomenal content. Since giving up this criterion would require that we forfeit our answers to a number of the preceding objections, we must block Shoemaker's argument for appearance properties.[6]

[6]Shoemaker recognizes that we do not seem to experience relational properties—especially not

More than one step in Shoemaker's argument can be questioned. First there is his claim that, in the hypothetical situation described, it is not the case that one group of individuals has non-veridical experiences while the other has veridical experiences. I don't see why there could not be such a situation. An experience is veridical when its content obtains.[7] Shoemaker seems to assume that experiences which carry the same information about the world have the same content (hence the same status as veridical or non-veridical in the same circumstances). This comes out when he asks "how can the experiences of Jack and Jill represent the tomato differently and yet neither of them misrepresent it, *given that the same information about its intrinsic nature is getting to both"* (my emphasis). I agree with Shoemaker that it is hard to see how the informational contents of Jack's and Jill's experiences could differ with regard to their accuracy given the same perceptual conditions. But virtualism is not about informational content. The contents virtualism is about are what one sees_{p}, hears_{p}, etc. There is nothing wrong with a scenario in which men see_{ip} blue bananas, women see_{ip} yellow bananas, and only women see correctly.

Even granting Shoemaker's first claim, we can still resist the argument by attributing non-veridical experiences to both groups of perceivers. On the face of it, there is nothing wrong with men $\text{seeing}_{\text{ip}}$ blue bananas and women $\text{seeing}_{\text{ip}}$ yellow bananas while bananas are red. Again, this would be hard to make sense of on the assumption that the contents of experiences are determined by what information they carry (or carry in normal conditions, or were designed to carry, etc.); Shoemaker's argument has force against virtualism combined with an informational account of relation R. However, the argument does not show that what one sees_{p} does not go hand-in-hand with the phenomenal character of one's visual experience.

appearance properties. He argues that this is an illusion due to the fact that we experience appearance properties *as* monadic properties. But there is arguably no room for one to see_{p} a property as something else than what it is, so Shoemaker's view does not seem to square with our position.

[7] Of course, this might not be what "veridical" means in everyday life since we give the term "content" a stipulative reading. Nevertheless, this is what it has to mean here for the thought experiment to be relevant.

6.8 Objections to projectivism

Shoemaker (1994) also suggests that the only viable alternative to the appearance property view is projectivism about color, the view that color experiences are always non-veridical. Why projectivism would be the only alternative to appearance properties is somewhat unclear. However, there are other considerations which favor projectivism about color, given virtualism and our notion of phenomenal content. In brief, there are no actually instantiated properties which might plausibly be taken to be what we see_p in color experience. In particular, we do not see_p physical-functional properties more than appearance properties. Our criterion of visual phenomenal content seems to require that we adopt projectivism about the contents of color experiences.

Projectivism is not exactly the orthodox view. Shoemaker, for one, rejects it. So something needs to be said in its defense. I cannot cover the topic in depth, but I want to indicate how I think projectivism should be defended. Let us look first at Shoemaker's reasons for rejecting it.

> This view has its own set of unattractive consequences. ... it implies that our perceptual experience is incurably infected with illusion—that we cannot help but perceive things as having properties that they do not and could not have. (1994)

So far, Shoemaker has merely restated the projectivist view in more dramatic language. What is so bad about systematically experiencing uninstantiated properties? Should not this be regarded as a genuinely open possibility given the nonfactive nature of experience?

> In addition, while we can make sense of the idea of there being properties that are in some way represented in our experience but never instantiated in anything—e.g., the property of being a ghost—it is difficult, to say the least, to make sense of the idea that experienced color could

> be such a property. Granted that there are in fact no ghosts, we at least have some idea of what would *count* as someone veridically perceiving an instantiation of the property of being a ghost. But if we ourselves do not count as veridically perceiving the instantiation of redness-as-we-experience-it, I think we have no notion of what could count as veridically perceiving this.[8] (1994)

Contrary to Shoemaker, I think I can easily see what it would be for an experience of red to be veridical even if all such experiences were non-veridical and nothing in the actual world were capable of making them veridical: it would be for something one is looking at to be red (by "red" here I mean the same thing one means by "red" when one talks about "experiences of red"). For one to have a veridical color experience is for the state of affairs one experiences in this experience (or a state of the kind associated with this experience, for derivative states) to obtain.

The real opposition to projectivism about color stems from the observation that we don't treat colors as mere projections in everyday life (Maund 1995; 2008). We routinely describe objects as colored. We order paint, clothes, and cars of particular colors, and we complain when we don't get the requested colors. Everyone agrees that grass is green and snow is white. The claim that nothing is really colored seems to fly in the face of common sense.

But the claim that things are not really the way they look to us is also now a piece of common sense. Everyone with a college education knows that objects are made of tiny particle-like things zipping about at high velocity in a near vacuum. And if there are no homogeneously filled surfaces, there are obviously no homogeneously red objects. Since color always appears in (apparent) homogeneous surfaces or volumes, it seems to be part of common sense that objects are not colored the way they look in experience. How are we to make sense of common sense?

[8] I can't resist pointing out that Shoemaker's choice of example is unfortunate: I only have the vaguest idea what it would take for there to be a ghost. Are ghosts necessarily transparent? Can there be ghosts which have never lived? Must ghosts be able to pass through walls? What is the difference between a ghost and a resurrected person?

It seems reasonable to suggest that color terms, like most terms of ordinary language, are used rather loosely in everyday life. Suppose you were to order a "blue" shirt from a store. If you were subsequently to accuse the storekeeper of not having sold you a shirt of the color you requested on the ground that nothing is really colored, and if she were sophisticated enough, she would say that every perceptible part of the shirt reflects the right wavelengths, and that this fulfills the order. Sometimes, we define or pick out colors in this way. But sometimes we don't. If I say that I saw_{ip} a blue dragon in my dream, I don't mean to suggest that the dragon reflected or emitted any kind of light, caused any kind of experience in me, or that it was made of any particular kind of material. By "blue", I mean a monadic property which is frequently presented to me in perceptual experience—what Chalmers (2006) calls "Edenic blue". Sometimes "blue" means Edenic blue, and sometimes it means something along the lines of *the property that normally causes experiences of Edenic blue in this world*.

There is a plausible explanation for why we are ambivalent like this. On the one hand, we have a strong tendency to recycle our words in the face of reference failure. Normally, when a term turns out not to refer on its usual understanding, we don't introduce better terminology to replace it; instead, we quietly change our use of the term to accommodate reality. We start using the term in a way that will insure that it refers and will save the apparent truth of the most important claims we want to make using it. On the other hand, we systematically believe the contents of our perceptual experiences: the default downstream causal profile of a perceptual experience is that of a judgment. It takes some effort not to fall in for what our perceptual experiences seemingly reveal to us. The perceptual beliefs about Edenic colors which are constantly pushed on us by the flow of experience are inconsistent with the scientific world view. Since the scientific view is also deeply ingrained, we end up with inconsistent beliefs or a tendency to wobble between inconsistent sets of beliefs. The result is that we sit on the fence regarding the use of color terms. If the

scientific view is right, our recycling policy recommends that we use them to refer to such properties as causing experiences of Edenic colors. If the phenomenal view is right, we can follow our initial inclination to use them to refer to the properties represented in color experience (Edenic colors). The dominant belief varies, and linguistic usage varies accordingly.

6.9 Summary

In this chapter I have tried to addressed a range of objections to propositionalism, one of the two central tenets of virtualism. There is a unifying thread in my responses to these objections: much of the initial plausibility of these objections has its origin in a conflation of the contents of the relevant experiences with other contents, e.g. the contents of propositional attitudes, the contents of emotions, the conventional meanings of public language expressions, or the informational contents of brain events. The objections tend to shed their plausibility once we focus on the contents which I claim satisfy virtualism: the proposition-like things we stand in the experiencing relation to, that is, the things we see_p, hear_p, feel_p, etc. I tried to show that these contents mirror phenomenal character.

Chapter 7

Disjunctivism and introspection

One of the main alternatives to a representational or virtual theory of experience is a disjunctive one. As noted in chapter 1, the core tenet of disjunctivism is that phenomenal states instantiated in veridical experience cannot be instantiated in non-veridical experience. The chief motivation for this claim is a kind of naive realism I dubbed *factualism*. According to factualists, the phenomenal states instantiated in veridical experience consist in standing in a factive relation to states of affairs.

The reason disjunctivists refer to factualism as "naive realism" is that they think it best captures the pre-theoretic conception of experience which is given to us through introspection. My aim in this chapter is to address this motivation for factualism. I will argue that the revelations of introspection militate for virtualism, not factualism. I will discuss other motivations for and against factualism (and disjunctivism) in the next chapter.

While factualism appears to be at the heart of most positions which go by the name of "disjunctivism", two potential exceptions were identified in chapter 1: it is possible that Martin (2004; 2006) and Snowdon (1980; 2005) do not endorse disjunctivism but a potentially weaker claim to the effect that experiences occurring in veridical and non-veridical conditions never share a "fundamental kind". Unless sameness of phenomenal character implies sameness of fundamental kind, this posi-

tion is compatible with virtualism. I set this interpretation aside because 1) it seems plausible that sameness of phenomenal character implies sameness of fundamental kind (whatever a fundamental kind is) and 2) my goal is to assess virtualism, not every theory which might have been taken as an alternative to virtualism. Snowdon's position might also not fall under the heading of the view I call "disjunctivism" if it is best characterized by the particularist thesis. Particularism is compatible with virtualism, so I will for the most part set it aside here. I will nevertheless address some of the evidence for it in the process of discussing the introspective motivation for factualism.

Factualists and virtualists agree that veridical experiences consist in standing in a relation to proposition-like things. Their disagreement lies in the fact that factualists think the relevant relation is factive while virtualists think it is virtual (hence non-factive). This makes it natural to speak of both views as ascribing contents to veridical experiences: the contents of veridical experiences are the proposition-like things we are related to in these experiences. That is how I will use the term "content" when talking about veridical experiences in this and the next chapter.

7.1 The introspective case for factualism

Fish (2009: ch. 1) offers one of the clearest statements of the introspective case for factualism. According to Fish, it is introspectively manifest that external objects and their properties "shape the contours" of the phenomenal characters of veridical experiences, and this means that these phenomenal characters are factive: they are states of standing in a factive relation to states of affairs. Martin (2004: 65; 2006: 354) and Hellie (2007) also explicitly claim introspective support for factualism (though they refer to it as "naive realism").

Fish and Hellie both offer the same collection of introspective reports to show that introspection supports factualism:

(7.1) Mature sensible experience (in general) presents itself as ... an immediate consciousness of the existence of things outside of us. (Strawson 1988: 97)

(7.2) The ripe tomato seems immediately present to me in experience. I am not in any way aware of any cognitive distance between me and the scene in front of me; ... The world is just there. (Levine 2006: 179)

(7.3) When someone has a fact made manifest to him, ... the obtaining of the fact is precisely not blankly external to his subjectivity. (McDowell 1982)

(7.4) In its purely phenomenological aspect seeing is ostensibly saltatory. It seems to leap the spatial gap between the percipient's body and a remote region of space. Then, again, it is ostensibly prehensive of the surfaces of distant bodies as coloured and extended ... It is a natural, if paradoxical, way of speaking to say that seeing seems to "bring one into direct contact with remote objects" and to reveal their shapes and colours. (Broad 1952)

(7.5) Visual phenomenology makes it for a subject as if a scene is simply presented. Veridical perception, illusion, and hallucination seem to place objects and their features directly before the mind. (Sturgeon 2000: 9)

Fish also cites Campbell and Martin on how their experiences strike them in introspection:

(7.6) The phenomenal character of your experience, as you look around the room, is constituted by the actual layout of the room itself: which particular objects are there, their intrinsic properties, such as color and shape, and how they are arranged in relation to one another and to you. (Campbell 2002: 116)

(7.7) Some of the objects of perception – the concrete individuals, their properties, the events these partake in – are constituents of the experience. No experience like this, no experience of fundamentally the same kind, could

have occurred had no appropriate candidate for awareness existed. (Martin 2004: 39)

While not all from salaried disjunctivists, the preceding quotes are a clear echo of disjunctivists' motto. Disjunctivists claim that the relation between phenomenal states and external states of affairs introspectively seem—both to them and to laypersons—to involve a "direct" or "immediate" contact, or some kind of "constitution", and that this observation militates strongly for factualism or disjunctivism.

7.2 The introspective case against factualism

Opponents of factualism are generally less explicit about the introspective support for their position. There is nevertheless an introspective case against factualism. My goal in this section is to expound this case. The argument ultimately rests on phenomenological premises, so I don't expect it to persuade any card-carrying factualist. This is not what I aim to do; for now, I only want to make a case against factualism which is apt to sway neutral parties. If I succeed at this, I will be in a position to deploy a more powerful argument against factualism in later sections.

To see how the phenomenological case against factualism goes, it is necessary first to be clear on what a virtualist can grant to the factualist regarding the deliverances of introspection.

When I introspect, I find "an immediate consciousness of the existence of things outside of" me. If I am looking at a tomato, "the tomato seems immediately present to me in experience." "I am not in any way aware of any cognitive distance between me and the scene in front of me." The state of affairs I am experiencing is "precisely not blankly external to [my] subjectivity." It is for me "as if a scene is simply presented." The facts "leap the spatial gap" between me and the world, and I come in "direct contact" with them—they are "directly before the mind". It seems undeniable that the scene I am experiencing is a "constituent" of my experience.

All of this is perfectly compatible with virtualism. My consciousness of external states of affairs is immediate as Strawson and Levine emphasize, but all this means is that I am not made conscious of external states of affairs through conscious inference or through awareness of other things. I am phenomenally conscious of the external state of affairs *tout court*, and without any conscious priming. The virtualist can also agree with Broad, McDowell, Sturgeon, Campbell and Martin that external states of affairs constitute or enter into our phenomenal states. Presumably, external states of affairs are states of affairs which involve properties of sorts only broadly-speaking material objects instantiate (e.g. shape properties). According to the virtualist, a phenomenal state *consists* in standing in a relation to a proposition, hence is constituted by a proposition. Since we count states of affairs as propositions, virtualism is compatible with the claim that phenomenal states are constituted by states of affairs, whether external or not. A virtualist can even agree that phenomenal states are constituted by *facts* in veridical cases: when a phenomenal state is constituted by a state of affairs which obtains, it is constituted by a fact, because facts are just states of affairs which obtain.[1] This is why I introduced the label "factualism". It seems to me that the quotes from Strawson, Levine, McDowell, Broad and Sturgeon are representative of important tenets of naive realism—a view I hold—and not factualism, the stronger view which motivates disjunctivism.

One point factualists commonly make which I have not granted is that ordinary objects enter into the natures of phenomenal states through the states of affairs we are related to in them (i.e. that particularism is true). I do not find this to be phenomenologically accurate. When I introspect my current computer monitor experience (in the everyday sense of "experience"), I do not find any trace of Eddy (my monitor). I cannot tell simply by scrutinizing my experience whether it is an experience of Eddy or an experience of Betty, a qualitatively indistinguishable

[1]This assumption has not gone unchallenged (see Mulligan 2008). There is no room to debate the nature of facts here. I will just say that if the virtues of factualism and disjunctivism turn on denying this, they have nothing to do with introspection.

monitor which has been delivered for someone else at the same time as Eddy. On this ground, I conclude that my phenomenal state is not specifically "of" Eddy (or Betty), though of course my experience is an experience of Eddy in the everyday, "encounter" sense of "experience": my encounter is with Eddy. Introspectively, I find that my phenomenal state relates me to a generic state of affairs involving no particulars. My phenomenal state tells me about a particular—it tells me that there is *a* monitor-like thing in front of me—, but it tells me this only on an unspecific reading: it tells me that there is a monitor-like thing in front of me.

The problem with particularism is most obvious when we consider distant objects we know nearly nothing about. Imagine for example that you are looking at a star in the night sky. You then close your eyes, spin on yourself, open your eyes again, and locate a star. You could easily end up looking at a different star than you were earlier while having exactly the same experience (while instantiating exactly the same phenomenal state). If the stars themselves were essential constituents of the phenomenal states you instantiate, these states would be different and we should expect the difference to show in introspection, thereby giving you a clue as to whether you are looking at the same star or not. You don't have a clue, so it seems that the stars you are looking at do not enter into the natures of your phenomenal states.

Of course, one can posit that there are phenomenal differences between experiences that are not introspectible. Insofar as this claim is supposed to apply to the sorts of ordinary experiences from which we derive our concept of a phenomenal state, it seems to me to be in tension with the very idea of a phenomenal state, but there is a more basic problem here. The preceding cases show among other things that our naive introspective judgments about our own phenomenal states do not reflect the external objects causing these states: they are not judgments in which we ascribe ourselves relations to external objects or object-involving SOAs. We are asking whether factualists' claim that naive introspection reveals phenomenal states

to be constituted in part by external particulars can motivate their view. Since our findings contradict factualists' starting point (their claim about what naive introspection reveals), they block cases of this kind for factualism. The postulate that experiences' involvement of external objects is not introspectible can save the view that we experience objects or object-involving SOAs, but it cannot restore the claim that the correctness of this view is introspectively evident.

The intuition that ordinary objects enter into the natures of phenomenal states can to a large extent be explained away by two facts.

First there is the fact, mentioned in introduction (p. 4) and repeated many times since, that experiences in the everyday sense are *encounters*. They are events akin to bumping into a wall. On *this* conception of experiences, most are unquestionably object-involving. If phenomenal states were the states exemplified by experiences in this sense, most would be object-involving. But this conception of experiences and phenomenal states trivializes the claim that certain phenomenal states are object-involving. It is not the conception which is relevant to this debate. I am talking about phenomenal states in a different sense introduced on page 4 (the felt components of sensory, emotional, and cognitive experiences in the everyday sense). I suspect that much of the appeal of particularism lies in the obviousness of a homonymous thesis about experience in the everyday sense.

A second factor is that perceptual experiences typically reveal enough information about external objects to allow us to uniquely identify them. Visual experiences, at least, generally present objects as existing at particular egocentric locations and belonging to certain general categories of object (e.g. material object, living creature, person). This information is generally sufficient to uniquely identify an object. Once one knows that an object of a certain general kind has been at a certain location at a certain time, one is normally in a position to refer to the object. That is so in virtue of the fact that (as a general rule at least) only one object of a kind can occupy any given location at a time. Empowered by the information provided

by perceptual experience, we systematically, seamlessly form thoughts with such contents as *the object at position x is F*. This transition from experiencing generic contents to forming definite thoughts about particulars is easily overlooked. One could easily come to be under the erroneous impression that perceptual contents themselves involve the equivalent of singular terms.

So I think there is good reason to deny that phenomenal states involve particulars in any interesting sense, and also a plausible explanation of why this might have seemed to be the case. Having said this, the virtualist *can* agree with Campbell and Martin that "the concrete individuals, their properties, the events these partake in" are "constituents of the experience." It is not part of the virtualist or representationalist position that the contents of experiences do not involve concrete individuals. It might seem that virtualism combined with the claim that individuals enter into the contents of experience would commit one to the view that there are non-existent individuals. But virtualism at most requires that SOAs can fail to obtain when we experience them. A singular SOA can fail to obtain while all the objects it involves exist: the properties predicated of the objects in the SOA can fail to be instantiated by the objects even if the objects exist. So virtualism does not imply that there are non-existent objects even on the assumption that all the SOAs we experience are singular. If a virtualist denies that experience has singular content, that is for reasons of phenomenologically accuracy. This is a matter that we can bracket for our purposes: if factualists are right about the singular character of the contents of experience, this is something virtualists can take on board.

The preceding remarks carry over to Fish's (2009: 22) claim that the SOAs we are related to in experience involve particular property instances. According to Fish, one does not merely experience a blue ocean or the fact that the Pacific Ocean is blue, but the Pacific Ocean's particular instance of blueness. Fish stipulates that by "the Pacific Ocean's particular instance of blueness" he does not mean a trope.[2]

[2]See his footnote 26.

It is not clear what he means in this case, but I am reasonably confident that I don't introspectively find the Pacific Ocean's particular instance of blueness as a constituent of any phenomenal state: if I did, I would be able to distinguish cases where I am looking at the Pacific Ocean from cases where I am looking at Twin Pacific Ocean simply by scrutinizing the phenomenal character of my experiences, but I cannot. So I don't think Fish's claim is phenomenologically accurate. Again, virtualists could also probably grant the point without committing themselves to Meinongian property instances. Take for example a state of affairs P which involves both the Pacific Ocean's instance of blueness and there being a golden mountain. Even if it is not possible to be related to the ocean's blueness without its obtaining (on pain of Meinongianism), it is clearly possible to be related to P while there is no golden mountain. So we can be related to non-obtaining states of affairs even if we always experience particulars of the kind Fish describes as property instances, whatever these are.

My point so far is that much of what has been said in support of factualism and disjunctivism is compatible with virtualism. In particular, the familiar observation that experience is immediate and constituted by external states of affairs—naive realism as traditionally understood, I think—does not the least support factualism, because it is part and parcel of the broader virtualist view.

A comparison with the case of belief should help drive the point home. I believe that Bob the mailman will pass by this morning. We can think of my belief as the state of standing in a certain propositional attitude relation to a certain state of affairs S. My belief that S is the case is direct and immediate: it is not a belief I have merely in virtue of having other beliefs, and it is not (need not be, in any case) the product of a conscious inference. One could also say that an external fact is a constituent of my belief, because my belief is a relation to S, and S is an external fact (it is an external state of affairs which obtains, as it happens). We can even allow that an individual (Bob) is part of state S. All of this is certainly compatible

with holding that the attitude relation which is constitutive of belief is virtual.

The preceding remarks are crucial to the virtualist's case from introspection against factualism. The claim that the virtualist should make about introspection is that *it does not present phenomenal states as having any property which is incompatible with their being instantiated independently of how the world is*—independently of the states of affairs presented in these experiences obtaining, at least. To use the language of factualists, introspection reveals an unmediated, constitutive openness of the mind toward external states of affairs, but (the virtualist should add) nothing more of interest. What introspection reveals is compatible with virtualism. This is momentous, because there is arguably nothing more to the essence of phenomenal states than what we can know about them through introspection. If there is no property which makes them factive and which we can know about through introspection, then they are arguably not factive. I will call this the *direct argument from introspection.*

The direct argument from introspection

1. If all veridical phenomenal states essentially have property F, then that some have F should be apparent in introspection.

2. Veridical phenomenal states do not introspectively seem to be factive.

   Therefore, veridical phenomenal states are not essentially factive.

Premise 1 seems highly plausible in light of how one comes to acquire the concept of a phenomenal state: one comes to acquire the concept of a phenomenal state by generalizing from examples of phenomenal states one can introspect. This leaves almost no room for veridical phenomenal states in general to essentially have properties which they cannot be seen to have in introspection. Premise 2 is justified in good part by the preceding discussion of what introspection reveals about phenomenal states. Introspection reveals relations to external states of affairs, perhaps even

relations to particulars, but this is compatible with virtualism. What introspection would have to reveal to falsify premise 2 is the modal property of factivity itself, and I deny that it reveals this. Factualists will disagree, but I am not yet trying to persuade them. At this stage I am merely trying to articulate the case from introspection for virtualism as it stands; we will see later how this case can be reinforced to yield a more persuasive argument.

My claim that there is nothing more to the essence of phenomenal states than what we can know about them through introspection could conceivably make a physicalist nervous: it looks like we would obtain a valid argument against physicalism by replacing premise 2 in the argument above with the claim that phenomenal states do not introspectively seem to be physical. But a modern physicalist ought to say that this latter claim is not trivial at all: when you know that your experience has a certain phenomenal character, it could well be that you thereby know that it has a certain physical or functional property, albeit under a mode of presentation which makes this opaque to you. This is the standard response to Jackson's parallel knowledge argument, so there is no more threat to physicalism here than there is in Jackson's claim that Mary does not know what it's like to experience red while she remains in her black-and-white room however much she learns about the physics of color vision. That is not to say that the knowledge argument is not sound, but that my premises do not go beyond what physicalists are already prepared to accept, by and large.[3]

It is worth noting that it is possible to make the case against factualism without

[3]One might ask why factualists cannot apply the same strategy. How do you know, one might ask, that what seems to you like a virtual relation is not a factive relation under a different mode of presentation? First, virtuality and factivity are incompatible. It is much harder to see how the story would go in this case. Moreover, all physicalists should agree that the introduction of modes of presentation is a somewhat ad hoc device. It is to physicalism what disjunctivism is to factualism: an accepted price to pay to save the theory. While physicalists can afford an ad hoc response due to the powerful independent motivations there are for physicalism, the situation is rather different with factualism. The factualist's position, if what I said about the revelations of introspection is correct, is something like that of a physicalist trying to turn Jackson's knowledge argument into an argument for physicalism. There is a tiny chance that we are being misled regarding the factive or nonfactive character of experience by the mode under which it is presented in introspection, but that is all there is in absence of independent evidence: a tiny, microscopic chance.

appealing to the potentially controversial premise 1:

The probabilistic argument from introspection

1. If a state does not seem to have property F essentially when inspected as well as we can, then it very probably does not have F essentially.

2. Phenomenal states do not introspectively seem to be factive.

3. Introspection is the best method we have of inspecting phenomenal states.

   Therefore, phenomenal states are very probably not factive.

The rationale behind the new premise 1 is that there are far more possibilities than there are necessities. For this reason, it would be irrational not to expect nearly all properties not to be essential of any given thing.

So far I have not talked about actual cases of illusion or hallucination. I have merely said that the concept of experience as it is given to us in introspection leaves open the possibility of phenomenal states occurring independently of corresponding facts, and that this is good reason to think that the possibility is genuine. We must not forget that there is also a more conventional argument from actual cases of illusion and hallucination which hinges mostly on introspection. The argument boils down to a simple observation: some actual illusory and hallucinatory experiences introspectively seem identical in phenomenal type to actual veridical experiences. If we are to believe introspection, we should conclude that they are in fact identical in phenomenal character, something that is ruled out by factualism.

Consider first the case of mirrors. While probably all mirrors have imperfections which result in overall visual experiences which are qualitatively different from veridical experiences in subtle ways, it seems likely that components of visual experiences caused by mirrors (an experience of a shape at a certain location, for example) are sometimes introspectively indistinguishable from actual veridical experiences. Experiences caused by mirrors generally represent objects as located

at positions behind where the mirrors would appear if they were represented in the experience. Since there are generally no relevant objects at these locations, experiences caused by mirrors are generally illusory. It also seems that they often match veridical experiences in phenomenal character as far as we can tell, so it seems that there are actual illusory experiences which introspectively appear to us exactly the same way as some actual veridical experiences introspectively appear to us. Alternatively, take the illusion of motion cinematography relies on. It is undeniable that displaced, discreet apparitions of a given shape at a sufficient speed can result in an experience which is introspectively identical to an experience of the same shape moving continuously along the same trajectory. The cinematographically induced experience is illusory because it presents continuous motion where there is none. If this case is unclear because it is unclear that experience presents motion, take instead this familiar illusion. All digital displays, including today's computer monitors, are made of pixel grids. There is always some separation between the pixels. But this separation is invisible from a normal viewing distance. The white background of my screen appears to be uniformly white, but it is not. We know that no object is densely filled (far from it), but my screen at least appears to have a level of uniform density real-world objects really have, but without having it. This is the second way that introspection speaks against factualism: veridical and non-veridical experiences can introspectively appear exactly the same way. Absent evidence to the contrary, we should believe what introspection tells us about our phenomenal states, so there is good reason to think that some phenomenal states can occur both in veridical and non-veridical conditions.

Let us take stock. Factualists claim that their view is supported by introspection, but this is also something that their opponents can plausibly claim. First, they can reasonably claim not to find in introspection any feature which would prevent phenomenal states from occurring independently of external facts. They can claim this without disagreeing with much of factualists' alleged introspective findings. In

particular, virtualists need not disagree with introspective reports to the effect that phenomenal states make us directly aware of external facts and are "constituted" by external facts, on a normal reading of these claims. Second, virtualists can plausibly say on the basis of introspection that some experiences which are known to be non-veridical are phenomenally identical to experiences which are veridical.

7.3 How to settle the dispute

As I indicated earlier, I know what most factualists will think in response to my arguments from introspection: they will think that they introspect better than me. But I think I introspect better than them, and I am not the only one to say the kinds of thing I said above. How are we to get beyond foot stomping?

Fish (2009) and Hellie (2007) propose that we survey the opinions of those who introspect best—philosophers. The product of their survey is the list of reports I gave in section 7.1. We have seen that most of these reports (all except Martin's) say nothing that is clearly incompatible with virtualism. This line of argument would be weak even if the reports supported factualism; as it is, I think we can set it aside. We can set aside surveys generally, because it is clear that the point of disagreement between factualists and virtualists is too subtle for most people to grasp on the spot. As Martin (2006) says, "when we come to state the differences between the two positions, we find ourselves talking in terms of notions of modality and constitution. One might be skeptical whether it could really be part of any common sense view that objects were or were not constituents of our experiences of them." As for the sophisticates who understand the distinction, there is little hope of eliciting theoretically unbiased judgments from them.

An oddity in the factualist position points to a means of resolving the debate. From the factualist point of view, introspection is highly misleading: factualists are committed to the claim that experiences appear factive whether we are per-

ceiving veridically or not. They posit that introspection is misleading in making non-veridical experiences appear phenomenally identical to veridical experiences, but why not posit instead that introspection is misleading in making all experiences appear factive to factualists? Given that others report that their experiences do not introspectively seem to them to be factive, this seems like a perfectly reasonable hypothesis. Of course, a parallel point also applies to virtualists, to some extent: given our disagreement about the deliverances of introspection, neither factualists nor virtualists should be confident in their ability to introspect.

At this stage, we should accept that introspection in some sense reveals different things to factualists and virtualists, and see if any position is more likely to be correct than the other despite this fact. We must stop taking our introspective judgments at face value. Instead, we should ask which position best explains the fact that we make these judgments. I don't mean to suggest that introspection does not involve some kind of infallible awareness of our phenomenal states. I am neutral on this. Whatever the case may be, there is room for error between the inner act of introspection (if any) and the ensuing judgment. However introspection works, the fact is that the product of factualists' apparently honest efforts at introspection is a judgment which is incompatible with their opponents' own apparent findings. In this situation, we ought not to take our apparent introspective findings at face value anymore. Rather than try to justify our theories on the basis of our diverging introspective judgments, we should look to introspective judgments as mere facts to be explained.

I am in effect proposing that we take the perspective of a third party who is unable to make introspective judgments of her own. Such an individual would have no prima facie reason to favor the putative introspective findings of proponents of one theory over those of their opponents. All she could do is ask which theory is most likely to be true in light of the fact that the two camps make the judgments they make.

The question, then, is which theory can best explain the apparent introspective findings of both factualists and virtualists, under all relevant conditions. The facts to be explained are:

A-virtuality When virtualists introspect in veridical conditions, they seem not to find factive states.[4]

A-factivity When factualists introspect in veridical conditions, they seem to find factive states.

Indistinguishability Everyone seems to introspectively find the same kind of state (factive or non-factive) in veridical and non-veridical conditions, and, for both virtualists and factualists, there are phenomenal states which introspectively seem the same in veridical and non-veridical conditions.

Since there are multiple competing theories of introspection and some of these theories are probably not entirely neutral regarding the present debate, we should understand what it is for one to introspectively seem to find that something is the case or not in a maximally non-committal way. This should be taken to mean that one forms some awareness of a state of affairs by using the primary means one normally uses to ascertain one's own mental states, in some broad sense of "awareness" which allows episodes of awareness which are not perception-like.

The probability calculus can help capture the dispute and the strategy I am proposing for resolving it. Let O be the totality of relevant introspective facts we can take for granted, i.e. the conjunction of A-virtuality, A-factivity, and Indistinguishability. Let V and F be the virtualist and factualist hypotheses, respectively. We can think of the probabilities of V and F independently of O as their prior probabilities, noted P(V) and P(F). We want to know what probability we should assign to each hypothesis given O. For hypothesis V, this is noted P(V|O). What we want to know, more specifically, is the ratio of P(V|O) over P(F|O): since it is reasonable

[4] "A-" stands for "apparent".

to assume that one of the two hypotheses is correct, we only want to know how much more (or less) plausible V is than F, given O. By Bayes' rule,[5] this ratio is equivalent to the following:

$$\frac{P(O|V)P(V)}{P(O|F)P(F)}$$

In the rest of this chapter, I will assess P(O|V) and P(O|F), that is, how well the virtualist and factualist hypotheses can explain the introspective data on hand. I will discuss P(V) and P(F) in the next chapter.

7.4 The virtualist's explanation

Let us start with the virtualist's explanation. Virtualists' explanation of A-virtuality is that they introspect correctly. Given virtualism, correct introspection should not reveal phenomenal states as factive. No difficulty here.

The virtualist also has a simple explanation of A-factivity (factualists' judgments about their veridical experiences). The key to the explanation is that it is easy to overlook the possibility that experience relates us to non-obtaining states of affairs. The possibility of virtual relations to SOAs is an idea which takes some time getting used to, and there is some evidence that factualists either consider it incoherent or have overlooked it. Take for example Fish's and Hellie's appeal to the testimonials listed in section 7.1. The main claim the testifying philosophers make which might seem to bear on factualism is that phenomenal states are constituted by external states of affairs. Fish and Hellie seem to believe that this claim militates for factualism, but it does not: phenomenal states could be states of standing in a *virtual* relation to external states of affairs, and in this case they would be constituted by

[5] Bayes' rule:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

external states of affairs without requiring that these states of affairs obtain. It seems that Fish and Hellie overlook the very possibility of virtual relations. Martin (2004, 2006) and Campbell (2002) provide another clue. They consistently describe the core tenet of their positions as the claim that experience is "relational", and they oppose this claim to representationalism. For example, Martin writes that "intentional theories of perception are committed to denying the relational nature of [veridical] experience, even if they are inclined to describe experience as if it were relational." (2006: footnote 7) Virtualists don't have to deny that experience is relational; there is no tension of the kind Martin implies within the representationalist view. Martin seems to neglect the possibility of virtual relations. In light of these observations, it does not seem far-fetched to suggest that factualists are subject to a subtle conceptual slip: they correctly observe that phenomenal states of the kinds which occur in veridical conditions are relations to and partly constituted by external states of affairs, and they wrongly infer from this that these states depend on external states of affairs obtaining because they neglect the possibility of virtual relations. Judging by what they say, in print and in conversation, what factualists see in introspection is what I granted above: perceptual experiences are constituted by external states of affairs, in that they are relations to external states of affairs. Because they neglect the possibility of virtual relations, factualists wrongly conclude that the veridical experiences we have, at least, require that the SOAs they relate us to obtain.

Of course, all of this depends on the cogency of the claim that one can stand in a relation to a state of affairs which does not obtain. One might doubt that this is possible. But it is part of the virtualist position I recommend that this is possible. The matter can be debated, but it is immaterial to the issue at hand. The question at hand is how likely A-factivity is *on the assumption that virtualism is true* (i.e. how high P(A-factivity|V) is). On the assumption that the kind of virtualism I recommend is true, it is possible to stand in a relation to a state of affairs which does not obtain, and it is easy to see how factualists might have overlooked this fact.

I said that what factualists see in introspection is merely that experiences are constituted in part by external states of affairs, a fact which is compatible with virtualism. But the virtualist can also allow that factualists in some sense "see" that experiences depend on external states of affairs when they introspect. The explanation remains the same: all there really is to see is that experiences are constituted by external states of affairs, but factualists' spontaneous introspective judgments go a little bit beyond this because these judgments are based on a possibly unconscious, perhaps automatic inference based on the incorrect assumption that relational properties require that their relata obtain in order to be instantiated.

There is a supplementary explanation of A-factivity a virtualist can appeal to: factualists confuse what their perceptual experiences tell them with what introspection tells them about these experiences. Perceptual experience normally engenders belief or at least an inclination to believe. Typically, someone who experiences an F will at least be inclined to believe that there is an F. We saw roughly how this works in chapter 2. It might be difficult to tell perceptually induced judgments from the deliverances of introspection. You find yourself thinking that a certain state of affairs obtains (there is an F), and you simultaneously find yourself thinking that your experience consists in a relation to this state of affairs. It is not hard to see how you might wrongly conclude that your experience consists in a factive relation to a fact, or that you have introspectively seen this.

There is yet another consideration which helps explain A-factivity on the assumption that virtualism is correct. I have repeatedly stressed that a sensory experience, in the everyday sense of "experience", is an encounter of a certain kind. Right now, I am experiencing a yellow post-it note on my computer, in the everyday sense of "experiencing". There is a perfectly intelligible sense in which no similar experiential event could have occurred had it not been for certain external, object-involving facts obtaining. My experience, qua encounter with a post-it note on a monitor, is essentially a factive relation to the fact that this very post-it note is stuck

on this very monitor (the relation of encountering a state of affairs is factive). If phenomenal states were the states instantiated in experiences in the everyday sense, factualism would be true. Now, most proponents of factualism use the word "experience" without defining it. This raises the possibility that factualists mistakenly extend a correct conclusion about experiences in the everyday sense to experiences in the technical sense. It could be that factualists introspectively find nonfactive experiences (technical sense) but do not recognize them as such because they are prone to think of them as experiences in the everyday sense.

Given the three preceding potential explanations—and that all three are mutually compatible—, a virtualist should not be surprised that some theorists end up saying the kinds of thing factualists say. The virtualist seems to have a good explanation of A-factivity.

The virtualist's explanation of Indistinguishability is a natural extension of her explanations of A-factivity and A-virtuality. According to virtualists, the same phenomenal states are present in veridical and non-veridical experiences. Virtualists find the same virtual states because they introspect correctly in both cases. As for factualists, they find the same factive states because they commit the same errors in the non-veridical case as in the veridical case: they assume that there are no virtual relations, they conflate perceptual and introspective judgments, and/or they conflate the ordinary and technical notions of experience.

7.5 The factualist's explanation

Like virtualists, factualists have a good explanation of their own apparent observations about veridical experiences, at least on the face of it: given factualism, it would seem likely that some theorists introspect correctly and find veridical phenomenal states to be factive. However, factualists have a harder time than virtualists explaining Indistinguishability and their opponents' apparent observations.

The virtualist explanations of A-factivity have no counterpart on the factualist side for A-virtuality. Regarding the first explanation, virtualists appear to be making a distinction which factualists don't make (the distinction between factive and non-factive relations to states of affairs). Whether factualists make this distinction or not, they cannot explain virtualists' putative error by saying that the latter are missing the distinction. It is also implausible that virtualists are making a distinction where there is none to be made. In general, it is much easier to overlook distinctions than to see distinctions where there are none to make. The second virtualist explanation of A-factivity also has no counterpart in the case of A-virtuality for factualists. In the case of veridical perception, a virtualist would normally agree that her experience is veridical. Yet she does not see her experience as factive. This cannot be explained by her conflating a perceptual judgment with an introspective judgment. The third possible explanation also has no counterpart. This is another case where virtualists make more distinctions—they distinguish between experiences in the everyday and technical senses while, the explanation posits, factualists fail to distinguish the two.

One explanation of A-virtuality a factualist can give is that it is the result of Indistinguishability (which would then have to be explained independently). One might say that the reason virtualists do not introspectively seem to find factive states is that their introspective judgments are influenced by their awareness of the fact that they cannot introspectively distinguish veridical from non-veridical states, knowledge which leads them to infer that phenomenal states are not factive. But how would phenomenal states appear to one before one draws such a conclusion? The factualist cannot allow that they appear non-factive to virtualists prior to their recognizing Indistinguishability, because this is the fact about introspection which is supposed to be explained by our recognition of Indistinguishability. However, if our default introspective view of phenomenal states had it that all are factive, the rational conclusion to draw in light of the fact that we cannot distinguish factive phenomenal states from other states would be that we sometimes seem to introspect

experiences when we don't have any. This would be the only conclusion to draw because the default introspective view of experience is definitional of the nature of experience. If we had reached the virtualist point of view through an inference from Indistinguishability, we should at least have asked ourselves whether we should save Indistinguishability by endorsing virtualism or an error theory of introspection. I don't recall myself asking such a question. On the face of it, all the factualist can say to explain A-virtuality is that virtualists fail to notice the factivity of experience despite of their best efforts; factualists have no explanation of exactly why this would be the case.

At first it might seem that the factualist can at least explain Indistinguishability. An obvious factualist explanation of Indistinguishability is that in both veridical and non-veridical cases the brain receives the same input (or a stimulus sufficiently far upstream). Since what one introspects is determined by the state of one's brain, one ends up introspecting alike in both cases. Since the phenomenal states which can occur as part of veridical experiences depend on external states of affairs but introspective states do not, it is not surprising that one finds oneself unable to tell directly in introspection whether one is undergoing a veridical experience or not.

While this explanation of Indistinguishability seems cogent at first glance, problems arise once we ask what introspection could be on the factualist view. What, exactly, is involved in it introspectively seeming to someone that such and such is the case? It seems plausible how one's mental states introspectively seem to oneself is either a matter of how one experiences one's own mental states or a matter of how one judges (or is inclined to judge) one's own mental states are. Either way, problems arise for the factualist explanation of Indistinguishability.

Suppose that introspective seemings are experiences. A factualist can say either that there are no non-veridical experiences at all, or that there are but that they token different phenomenal states than veridical experiences. The first option is ruled out if introspective seemings are experiences, because the data show that we sometimes

introspect incorrectly. Indistinguishability and A-factivity, in particular, together entail that it sometimes introspectively seems to factualists that they instantiate factive states when they do not. The second option leaves two broad possibilities open: either non-veridical experiences token states of standing in nonfactive relations to certain entities, or they token phenomenal states which do not have a relational character. Neither of these options is attractive from the factualist standpoint. The first, that non-veridical experiences are nonfactive relations, makes Indistinguishability highly mysterious. Why should we find the nonfactive relational states tokened in non-veridical perception indistinguishable from the factive states tokened in veridical perception? To explain this, factualists would at least have to give an account of how the nonfactive, non-veridical experiences we are capable of acquire their contents. Without an account of the contents of non-veridical experiences and how they come about, we have no reason to expect them to match veridical experiences in a way that would explain Indistinguishability. The situation is even worse if we posit that non-veridical experiences are not relational. Then it is entirely mysterious own introspective experiences could reveal anything about mental states.

Suppose now that how one's mental states introspectively seem to one is a matter of how one judges or is inclined to judge one's mental states are. Consider a case of known illusion. Say for example that you are looking at a stick half immersed in perfectly transparent, completely invisible water. Knowing that the stick is half immersed in water, you recognize the illusion for what it is. Now imagine that a bent stick is held above the water. The stick is bent and displayed in such a way that it looks exactly like the straight stick. So you are having two indistinguishable experiences, one of which is veridical and one of which is not. This would seem to be just the sort of case that one must explain to explain Indistinguishability. But if introspection were a kind of judgment, we would expect it to be influenced by what ones believes. So we would expect your knowledge of the illusion to prevent you from judging that your experience caused by the first stick is an experience in which you

stand in a factive relation to a bent stick. Provided that you are sufficiently aware of the illusion, you should not even be inclined to judge that the stick in water is bent. On the conception of introspection we are considering, it should not introspectively seem to you that you are (factually) related to a bent stick, as far as your illusory experience of the bottom stick goes. This contradicts Indistinguishability, which requires that we introspectively find the same kind of state whether in veridical or non-veridical conditions.

I have argued that a factualist explanation of Indistinguishability is elusive whether introspection is thought of as a kind of judgment or a kind of experience. There are of course multiple ways out of this conundrum. A factualist could say that acts of introspection are cognitively insulated judgments, experiences which make exception to factualism, or sui generis mental events. The nature of introspection is obscure enough, there is plenty of wiggle room for factualists to escape the dilemma. Still, I think we can say with some confidence that it is much less clear how to explain Indistinguishability on the factualist view than it is on the virtualist view. On the latter, any account of introspection or introspective seemings goes. On the former, a special conception of introspection and non-veridical experiences appears required, and it is far from clear at this stage whether a suitable account could be provided.

7.6 Summary

We have seen that both factualists and virtualists can claim some prima facie support from introspection for their views. Since both parties seem to introspect the best they can, I have suggested that we should not take their apparent introspective findings at face value. The only relevant findings that we are entitled to take for granted are A-virtuality, A-factivity, and Indistinguishability, which are facts about what we *seem* to find when we introspect. I have then argued that virtualism offers a far more plausible explanation of these facts than factualism. Put in terms

of the formalism introduced in section 7.3, I have argued that P(O|V) is far greater than P(O|F). If we were to give equal values to P(V) and P(F), the probabilities of virtualism and factualism independently of considerations having to do with the revelations of introspection, we would be in a position to conclude that virtualism is far more plausible than factualism. In the next chapter, I will argue that we should also assign a significantly greater credence to virtualism than to factualism independently of the revelations of introspection.

Chapter 8

Being in contact with the world

In the preceding chapter I argued that virtualism provides a better explanation of the facts about introspection everyone can agree on than factualism. It remains to see whether other considerations could sway the balance of evidence in favor of factualism.

Factualists often motivate their view by reference to its putative ability to secure a type of contact with the external world which representationalism does not allow (c.f. Campbell 2002; Child 1994b; 1994a: 145-9; Fish 2009: 23-4; McDowell 1982; 1986; 1994; 1995; 2002; 2008). Two potentially independent lines of argument can be discerned within this theme. First there is the epistemological case for factualism, which appears to have originated with McDowell. The idea is roughly that factualism has the merit of blocking at least some arguments for skepticism about the external world. There is also a potentially distinct line of argument to the effect that factualism is the only position which can account for the fact that our experiences and beliefs are "about" the external world at all (also apparently due to McDowell). Both arguments will be discussed in this chapter. I will end by raising a problem of my own for factualists concerning our ability to enter in contact with the external world.

8.1 The epistemological argument

I will first follow McDowell's exposition of the epistemological argument in his 1995 article "Knowledge and the internal". Then I will look at an alternate formulation due to William Fish (2009).

McDowell sets out the argument as a reductio ad absurdum: if certain assumption*s* are accepted, skepticism follows; since skepticism is false, one of these assumptions must be rejected, and the best candidate is the assumption that "standings in the space of reasons" are internal. Rejecting this assumption supposedly leads to endorsing factualism (McDowell skips factualism and concludes to disjunctivism, but we will see that the argument implicitly goes through factualism). We will see shortly what McDowell might mean by "standings in the space of reasons". This expression is best taken as a placeholder to start.

McDowell begins by stating a point he apparently takes to be foundational:

> I am going to work with an idea from Sellars, that knowledge–at least as enjoyed by rational animals–is a certain sort of standing in the space of reasons. (1995: 877)

He then states the view he aims to criticize:

> My concern is a familiar philosophical dialectic, which I shall approach in terms of what happens to the Sellarsian idea when the image of standings in the space of reasons undergoes a certain deformation. [...] The deformation is an interiorization of the space of reasons, a withdrawal of it from the external world. (ibid.)

The view that is McDowell's target is that standings in the space of reasons are internal. McDowell does not define "internal". I will articulate his argument formally before trying out various interpretations of its key terms.

According to McDowell, problems arise for the view that standings in the space of reasons are internal when we consider illusions and hallucinations:

> Consider the Argument from Illusion. Seeing, or perhaps having seen, that things are thus and so would be an epistemically satisfactory standing in the space of reasons. But when I see that things are thus and so, I take it that things are thus and so on the basis of having it look to me as if things are thus and so. And it can look to me as if things are thus and so when they are not: appearances do not give me the resources to ensure that I take things to be thus and so on the basis of appearances only when things are indeed thus and so. [...] So if I want to restrict myself to standings in the space of reasons whose flawlessness I can ensure without external help, I must go no further than taking it that it looks to me as if things are thus and so. One might hope that this inward retreat is only temporary. [...] The hope is that I might start from the fact that things look that way to me; add in anything else that the ground rules allow me to avail myself of, if it helps; and move from there, by my own unaided resources, without needing the world to do me any favors, to a satisfactory standing in the space of reasons with respect to the fact that the world is arranged the way it looks. [...] Anyone who knows the dreary history of epistemology knows that this hope is rather faint. (1995: 877-8)

The argument McDowell opposes starts with the observation that internal standings in the space of reasons cannot be factive states such as knowings of external states of affairs. Since knowings are standings in the space of reasons, this forces us to conclude that knowledge of external facts is impossible (on the assumption that standings in the space of reasons are internal). The argument may be schematized as follows:

Argument S

1. States of knowing are standings in the space of reasons.

2. Standings in the space of reasons are internal.

3. If standings in the space of reasons are internal, they do not entail external facts.

4. The state of knowing a fact entails this fact.

   Therefore, knowledge of external facts is impossible.

To safeguard the possibility of knowledge of external facts, we must reject one of the premises of this valid argument. McDowell's own view is that we should reject premise S2. It is the rejection of S2 which is supposed to motivate factualism and disjunctivism.

Given that premises S3 and S4 seem fairly secure, we can collapse them into one premise for simplicity. We can also rewrite McDowell's argument as a straight inference rather than a reductio:

Argument N

1. States of knowing are standings in the space of reasons.

2. If standings in the space of reasons are internal, they are not states of knowing external facts.

3. Knowledge of external facts is possible.

   Therefore, some standings in the space of reasons are not internal.

The argument is incomplete as it stands because it does not say anything about phenomenal consciousness (that is, it is incomplete qua argument for factualism). To take us to factualism or disjunctivism, the argument needs an additional premise to the effect that—as McDowell might put it—one's standing in the space of reasons is internal to one's subjectivity:

> [O]ne's epistemic standing on some question cannot intelligibly be constituted, even in part, by matters blankly external to how it is with one subjectively. For how could such matters be other than beyond one's ken? And how could matters beyond one's ken make any difference to one's epistemic standing? [..] the disjunctive conception of appearances shows a way to detach this 'internalist' intuition from the requirement of non-question-begging demonstration. When someone has a fact made manifest to him, the obtaining of the fact contributes to his epistemic standing on the question. But the obtaining of the fact is precisely not blankly external to his subjectivity [..] (1982)

I take it that the subjectivity in question is phenomenal consciousness; otherwise, it is hard to see how we might eventually reach a conclusion which bears on factualism or disjunctivism. Plausibly, that standings in the space of reasons are internal to phenomenal consciousness would require at least that they supervene on phenomenal states (past, present, and future). If we add this claim as a premise we obtain an argument which takes us close to factualism:

Argument M

1. States of knowing are standings in the space of reasons.
2. If standings in the space of reasons are internal, they are not states of knowing external facts
3. Perceptual knowledge of external facts is possible.
4. Some standings in the space of reasons are not internal. (From 1-3)
5. Standings in the space of reasons supervene on phenomenal states.
6. If some standings in the space of reasons are not internal and standings in the space of reasons supervene on phenomenal states, then some phenomenal

states are not internal.

Therefore, some phenomenal states are not internal.

So far, we have left the terms "internal" and "external" undefined. Independently of what these terms normally mean, we can stipulate that a state is external (i.e. not internal) just in case it is factive. Only this definition, or a stronger definition, yields a conclusion which bears on factualism. Read this way, the conclusion of argument M guarantees that at least some phenomenal states confirm factualism. It is only a matter of finessing to extend this argument to a large proportion of veridical phenomenal states, so we can grant that the conclusion warrants factualism.

A residual difficulty is that it is not clear what standings in the space of reasons are. McDowell does not really explain this expression. There is no explanation in his 1995 article, and all I could find elsewhere is a quote of this passage from Sellars' *Empiricism and the Philosophy of Mind* in *Mind and World* (p. xiv):

> The essential point is that in characterizing an episode or a state as that of knowing, we are not giving an empirical description of that episode or state; we are placing it in the logical space of reasons, of justifying and being able to justify what one says. (Sellars 1956/1997: 8)

This is the only place where Sellars uses this expression in EPM. Judging from this passage, having a standing in the space of reasons would seem to be a matter of being able to justify certain claims, but it is not clear whether Sellars (or McDowell) intends the specific manners in which one can justify a claim to be part of one's standing in the space of reasons. Take for example the state of being able to justify the belief that P based on one's *knowledge* that P is the case. Does this count as a standing in the space of reasons?

We can get around this terminological difficulty by assessing the argument on the two relevant kinds of reading of "standing in the space of reasons". Standings in the space of reasons can be such that all knowings are standings in the space

of reasons, or they can be such that not all knowings are standings in the space of reasons. We can take the standings referred to by McDowell to be standings of either kind. However, the argument only has a chance on the first reading, because the second makes M1 trivially false, so we have to take McDowell to be using "standing in the space of reasons" in a sense in which all knowings are standings in the space of reasons.

The problem with the argument on this reading is that M5 seems highly dubious. If states of knowing external facts count as standings in the space of reasons, M5 entails that external knowledge supervenes on phenomenology, which is highly dubious. Take hallucinations. If there are phenomenally identical pairs of hallucinations and veridical experiences, then it is almost guaranteed that two individuals' phenomenal profiles throughout their lives could be identical while one knows that P and one does not. Of course, disjunctivists such as McDowell deny that there are phenomenally identical pairs of hallucinations and veridical experiences. Still, the argument appears question begging on the present reading, because premise M5 has no appeal independently of disjunctivism on this reading. M5 seems dubious even assuming factualism and disjunctivism. The reason is that we have a certain amount of freedom in forming beliefs based on experience. Someone who is more prudent than most might fail to form the belief that P (hence fail to know that P) where another who has the same phenomenal profile would believe and know that P.

Let us consider an alternative statement of the epistemological argument due to Fish (2009) (the quotes are from McDowell):

> [Skepticism] gets a foothold only if we have to view perceptual experience in such a way that "even if we focus on the best possible case, [a subject's] experience could be just as it is, in all respects, even if there were no red cube in front of her". It is against the background of such a view of experience that the skeptic can make a compelling case

> that empirical knowledge is thereby shown to be impossible and hence that we do not and cannot have knowledge of the external world. ... it would constitute a response to the skeptic if we could show that we can "make sense of the idea of direct perceptual access to objective facts about the environment" (2009: 24)

As before, the motivation for factualism is that skepticism follows if factualism is rejected. Fish does not say exactly how the skeptical argument factualism addresses goes, but we can infer from what he says that the argument has the following top-level structure:

Argument S2

1. A subject's experience could be just as it is, in all respects, even if the subject were suffering from massive hallucination.

2. If (1), then we cannot have empirical knowledge.

   Therefore, we cannot have empirical knowledge.

So Fish's version of the epistemological argument for factualism, put in the form of a straight argument rather than a reductio, is:

Argument F

1. We can have empirical knowledge.

2. If (1), then a subject's experience could not be just as it is, in all respects, if the subject were suffering from massive hallucination.

3. A subject's experience could not be just as it is, in all respects, if the subject were suffering from massive hallucination. (from 1 & 2)

4. If (3), then factualism is true.

Therefore, factualism is true.

Once the argument is put this way, it loses all intuitive appeal it might have had. For why accept premise F2? Premise F2 simply states that the possibility of empirical knowledge requires the truth of disjunctivism (or some variant of disjunctivism). Fish does not say why we should believe this claim, and it is far from prima facie plausible. Suppose for example that someone tells you that it is raining outdoors, and that you have every reason to believe them. As a result, you form the belief that it is raining outdoors. According to premise F2, you could not count as knowing that it is raining unless the phenomenal states which are veridical in this world were such that they could not possibly have been non-veridical. This seems prima facie implausible.

Fish's and McDowell's cases suffer from parallel problems. The general form of argument is that skepticism threatens unless factualism is true. The key question is whether there really is a plausible argument for the conclusion that empirical knowledge is impossible which one can avail oneself to *on the assumption that disjunctivism is false*. Otherwise, anti-skepticism does not stand to benefit from factualism or disjunctivism. McDowell and Fish have not provided or uncovered such an argument. On the reading of "standing in the space of reasons" which ties such states to phenomenology, the skeptical argument rebutted by McDowell (argument S) assumes a relation between knowledge and phenomenology which is not acceptable under the supposition that disjunctivism is false (the relation captured by premise S1). Similarly, the second premise of the argument Fish responds to (argument S2) essentially states that the possibility of empirical knowledge implies disjunctivism, which is implausible on the assumption that disjunctivism is false. While disjunctivism does undermine the skeptical arguments addressed by McDowell and Fish, so does anti-disjunctivism.

8.2 The darkness-within objection

A common theme throughout McDowell's and other disjunctivists' work is that mental states—experiences, in particular—can only have content if they are constituted by these contents, in a way that makes them dependent on their contents obtaining. Otherwise, everything is dark inside. I will refer to this claim as the *Extension Requirement*. Sometimes McDowell puts the point as a rejection of "the dualism of conceptual schemes and content".

McDowell's arguments for the Extension Requirement are well hidden. I have searched his entire body of work on perception and found only two points which look like reasons for endorsing this claim.

The first point occurs in the following passage, which follows McDowell's introduction of disjunctivism in his 1982 article. He refers to disjunctivism as "the innocent position":

> We arrive at the fully Cartesian picture with the idea that there are no facts about the inner realm besides what is infallibly accessible to the newly recognized capacity to acquire knowledge [introspectively]. What figures in the innocent position I have just outlined as the difference between the two disjuncts cannot now be a difference between two ways things might be in the inner realm, with knowledge of which is the case available, if at all, only with the fallibility that attends our ability to achieve knowledge of the associated outer circumstance. Such differences must now be wholly located in the outer realm; they must reside in facts external to a state of affairs that is common to the two disjuncts and exhausts the relevant region of the inner realm. We cannot now see the inner and outer realms as interpenetrating; the correlate of this picture of our access to the inner is that subjectivity is confined to a tract of reality whose layout would be exactly as it is however things stood out-

> side it, and the common-sense notion of a vantage point on the external world is now fundamentally problematic. (McDowell 1982/1998: 241)

McDowell's discussion of the Cartesian picture is tangential to his justification of the Extension Requirement, because his reasoning would apply to any view which has relevantly similar implications. The relevant implication of the Cartesian picture is that phenomenal states are independent of the external world. On this picture, they must be independent of the external world because we have infallible access to them but not to the external world. According to McDowell, this implies that phenomenal states (subjectivity) and the external world do not "interpenetrate", and this in turn means that phenomenal states lack content about the external world.

This reasoning is easier to follow if we "picture subjectivity as a region of reality" (1982/1998: 240-1) as McDowell suggests we do. Think of phenomenal states as states in which we catch objects and states of affairs using a "phenomenal lasso". Phenomenal lassos are just like normal lassos, except that they are perfectly transparent and constitute experiences. The content of your experience at a time is what you have caught in your phenomenal lasso at that time, and a phenomenal state is a state of having caught certain entities in one's phenomenal lasso. On this picture, we can only sustain the view that phenomenal states are independent of the world by supposing that they are catchings of internal entities. I believe that McDowell endorses the Extension Requirement because he is thinking of phenomenal states along these lines.

But that is evidently not how experience works. Or so do virtualists claim. The factualist is not entitled to assume that anything like the lasso picture of experience is correct at the outset of an argument for factualism or disjunctivism, because that picture is flatly incompatible with the virtualist view of experience. McDowell's argument presupposes an uncommon picture of experience which is only appealing from within the factualist doctrine.

The other defense of the Extension Requirement I have found is in McDowell's

1995:

> If the space of reasons as we find it is withdrawn from the objective world as it makes itself manifest to us, then it becomes unintelligible how it can contain appearances, content-involving as they must be, either. We are here in the vicinity of a third Sellarsian idea, that reality is prior, in the order of understanding, to appearance; I am drawing the moral that it makes no sense to suppose that a space sufficiently interiorized to be insulated from specific manifest facts might nevertheless contain appearances. (1995: 889)

McDowell appears to be referring to Sellars' discussion of the "logic" of look statements. Sellars does make a claim to the effect that "reality is prior, in the order of understanding, to appearance". However, factualism cannot be inferred from Sellars' position. Sellars' point about appearances is summarized in this passage of EPM:

> The point I wish to stress at this time, [...] , is that the concept of looking green, the ability to recognize that something looks green, presupposes the concept of being green, and that the latter concept involves the ability to tell what colors objects have by looking at them [...] (1956/1997: 43)

Sellars' point, generalized, is that one cannot obtain the concept of something looking F prior to having the concept of something being F. This is the trust of Sellars' discussion of looks, which is primarily aimed at repudiating the logical atomist idea that concepts of external things can be reconstructed from concepts of observable facts which can be known with certainty. Factualism does not begin to follow from this. Virtualism, for example, analyzes experiences (appearances) as virtual relations to state of affairs such as something being an F (as opposed to looking F), and this very strongly suggests that reality is prior to appearance in the order of

understanding, yet virtualism obviously does not entail factualism. In other words, virtualism agrees with Sellars but not factualism.

Bill Child (1994a) develops another argument which he claims to find in McDowell (the "line of thought" referred to is McDowell's Extension Requirement):

> According to that line of thought, one cannot have thoughts about Fs unless: either one is (or has been) in direct cognitive contact with Fs; or one can construct a way of thinking of Fs from concepts of kinds of thing with which one is (or has been) in direct cognitive contact. Now on the non-disjunctive conception of experience we are not in direct cognitive contact with the world, since the most basic mental characterization of experience is world-independent. But it is arguable that no concept constructed solely from world-independent contents can itself be a concept of an objective world independent of thought. (147)

The following passage sheds some light on the reasoning behind Child's claim that we can be in "direct cognitive contact with the world" only if disjunctivism is true:

> [...] to think of conscious experience as a highest common factor of vision and hallucination is to think of experiences as states of a type whose intrinsic mental features are world-independent; an intrinsic, or basic characterization of a state of awareness will make no reference to anything external to the subject. (146)

Child does not seem to have a particularly demanding understanding of "cognitive contact", but he seems to have a very demanding understanding of the requirements of the highest common factor view (of the denial of disjunctivism): according to him, this view does not allow that an "intrinsic" or "basic" characterization of a phenomenal state includes expressions which refer to external entities, e.g. terms like "square" or "red".

Child's assumptions about the implications of the highest common factor view seem implausible in light of an analogy with the case of belief. Almost everyone (disjunctivists included) agrees that beliefs satisfy a highest common factor view: true beliefs and false beliefs are exactly the same qua beliefs; truth and falsehood are contingent properties of beliefs. This does not incline us to say that intrinsic or basic characterizations of beliefs cannot involve terms which refer to external properties or states of affairs. Consider in particular two of the main contenders for a basic characterization of the nature of belief: wide functionalism and interpretationism. On either view, the most basic characterization of a belief state might well make reference to external properties and states of affairs. In the first case, it could make reference to external properties or states of affairs as part of the functional role of the belief. In the second case, it could turn out that the most basic characterization of a belief state is along the lines of *being interpretable as believing that P*, and "P" could well be a sentence which involves such terms as "square" and "red" or picks out an external state of affairs. If the highest common factor view about belief allows terms which make reference to external properties as part of basic characterizations of beliefs, why not the highest common factor view about experience?

There are three main kinds of basic characterization of the nature of experience one might reasonably expect within the virtualist framework: narrow physicalist and functionalist accounts, wide physicalist and functionalist accounts, and dualist accounts. If dualism is true, "standing in R to $\exists x\, red(x)$" (or "standing the experiencing relation to $\exists x\, red(x)$") is the most basic characterization of an experience of red we can expect, and it makes reference to an external property. If wide physicalism or functionalism is true, again, the most basic characterization of an experience of red may be expected to make reference to redness (as in the case of belief). It is only on a narrow physicalist or functionalist account that it seems unlikely that the most basic characterization of experiences of red would make reference to redness. Child's line of reasoning at best yields an argument against narrow physicalism and

functionalism, not an argument against the highest common factor view.

Campbell (2002) also puts forward a related argument. He is primarily arguing against representationalism, but he also presents the argument as a case for the Extension Requirement, because he assumes that the alternative to representationalism is a view whose central tenet is the Extension Requirement. The argument turns on an appeal to the explanatory role of experience:

> Experience is what explains our grasp of the concepts of objects. But if you think of experience as intentional, as merely one among many ways of grasping thoughts, you cannot allow it this explanatory role. Suppose someone said: "Actually, reading newspapers is the fundamental way in which you understand the concepts of a mind-independent world. All your conceptual skills depend on your ability to read newspapers." The natural response to this would be that reading newspapers does indeed involve the exercise of conceptual skills, but it is simply one way among many of exercising those conceptual skills. Just so, if all there is to experience of objects is the grasping of demonstrative thoughts about them, then experience of objects is just one among many ways in which you can exercise your conceptual skills. At this point we do not have any way of explaining why there should be anything fundamental to our grasp of concepts about experience of objects. (122)

The argument seems to go something like this:

1. Experience is more basic than conceptual representation.

2. If representationalism is true, then experience is not more basic than conceptual representation.

   Therefore, representationalism is false.

Premise 2 is false if virtualism counts as a kind of representationalism. Virtualism says nothing at all about the conceptual or non-conceptual status of experience, or

its relation to propositional attitudes and concepts. It leaves it entirely open that all conceptual representations (on any reasonable understanding of this expression) derive from and depend on relation R in some way or other (hence that experience is more basic than conceptual representation). If conceptual representation is the kind of thing that is going on in propositional attitudes, this is essentially the view I gestured toward in sections 2.4 and 6.5. I would be very surprised if this view were inconsistent with virtualism.

Campbell's general approach to these matters in *Reference and Consciousness* is somewhat puzzling. He begins by arguing that conscious attention provides "knowledge of reference of demonstratives". The latter notion is explained as follows:

> Knowledge of reference of the demonstrative is what causes and justifies the use of particular procedures to verify and find the implications of propositions containing the demonstrative. (25)

The "explanatory role of experience" Campbell refers to in the first quote above when he objects to representationalism seems to be the role of conscious attention in providing knowledge of reference of demonstratives. Campbell's main point against representationalism is that only the sort of factualist view of experience he recommends can account for the explanatory role of experience. But the explanation of knowledge of reference in terms of conscious attention Campbell himself provides at the beginning of the book does not seem to depend on the factualist view. His explanation is that conscious attention directs sub-personal processes to construct *representations* of objects at particular perceived locations:

> To sum up: if you are to act intentionally on an object, you must consciously attend to it, in the common-sense use of the term; but that act of attention must also cause the selection of suitable information for processing, and suitable processes to operate on it, if the information-processing of which you are capable really is to be harnessed to your objectives. (27)

On the face of it, a virtualist could account for the role of conscious attention in just the same way as Campbell.[1]

8.3 The curse of the senses

According to factualists, the phenomenal states which are tokened in veridical experience are states of acquaintance with facts. One problem with this view is that, for most phenomenal states, there does not seem to be actually obtaining states of affairs which can account for their phenomenology.

Let us begin visual phenomenal states. As a general rule, in vision we experience the world as composed of densely filled objects. By "densely filled", I mean that the objects we experience seem to occupy all the space they span. But according to current scientific theory, there are no densely filled medium-sized objects. Rather, there are clouds of high velocity particle-like things whose interactions produce an illusion of densely filled objects. That is not to say that there are no medium-sized objects such as chairs and tables, but that if there are, they are not at all the way they appear to us in experience: they are something like four-dimensional clouds of particles. It seems obvious that the states of affairs we experience have no match in this picture of the world. It seems prima facie implausible to say that an experience of a red ball is a relation to a cloud of particles reflecting such and such electromagnetic wavelengths. One could say that this is what it is in spite of appearances—that the "mode of presentation" of experiences makes the identity opaque to us, but why should we believe this given the availability of the virtualist position? There is a strong prima facie case from science that visual experience is always illusory with respect to colors and shapes, so far as colors and shapes are supposed to be what we are related to in color and shape experiences.[2] Since all visual experiences have

[1] I think that is also true of Brewer's (1999) account of demonstrative reference, but I don't have space to engage with Brewer's reasoning here.

[2] The qualification that colors should be understood as the properties we experience in color experience is important. There are accounts of color which are plainly compatible with the scientific

colors and shapes as part of their contents, this means that all visual experiences are partially illusory. The same goes for other modalities: experiences of sounds are not relations to vibrations in a medium, taste experiences are not relations to chemical compounds, bodily experiences are not relations to any kind of physiological activity, and so on.

The moral we should draw from this depends on whether we read factualism as a claim about all possible veridical experiences or a claim about actual veridical experiences only.

Weak factualism The phenomenal states which actually are tokened in veridical perception are states of standing in a factive relation to external states of affairs.

Strong factualism The phenomenal states which can be tokened in veridical perception are states of standing in a factive relation to external states of affairs.

So far, I have been assuming the strong reading because I take factualists to be making a claim about the nature of experience generally, but the scientific picture of the world could conceivably lead one to fall back on the weak reading.

Call WF- the view that weak factualism is true but strong factualism is not. If WF- is correct, there are uninstantiated phenomenal states which could be both veridical and non-veridical (in different circumstances), yet the phenomenal states which are tokened in veridical perception in the actual world are necessarily veridical. One major oddity with WF- is that its truth seems to be up to us. Suppose that it has so far been confirmed by all actual phenomenal states. Presumably, some of

picture. For example, there is the view that colors are reflectance properties. It might be that this is what colors are in the everyday sense of "color". Personally, I find the everyday concept too muddled to lend itself to any clear analysis, as I explained in section 6.8. But we can set this aside. The point is only that no instantiated properties are plausible candidates for the properties which are characteristic of the SOAs we are related to in "color experience". If we call "colors" the properties which are characteristic of the SOAs we are related to in color experience, then colors are plausibly not instantiated. Whether these properties are what we mean in everyday life by "colors" is immaterial to the issue at hand.

Chalmers (2006) pursues the implications of the scientific view of the world at much greater length than I can afford to here.

the contingently possible phenomenal states which contradict strong factualism are states we could in practice bring about (thereby falsifying weak factualism). So it seems to be up to us whether WF- will continue to be confirmed. The position is literally unstable. Independently of this, if the actual veridical phenomenal states are factive, we should expect that this is a feature of the general kind of phenomenal state they fall under (a kind of state which has many uninstantiated instances). We should expect factualism to be true on the strong reading if it is true on the weak reading.

Factualists also seem committed to the stronger reading by their argument from introspection (without which the view would lose much of its appeal). According to factualists, naive introspection presents us with factive states and nothing else. Factualists are right to insist that this putative datum militates for their view (what I denied earlier is that it is a datum). We all agree that introspection yields judgments of the form *I am in state S*, and that phenomenal states are states of the kind exemplified by the states we self-ascribe in this way. If all the relevant states we self-ascribe in this way were factive as factualists claim, phenomenal states would be by nature factive, because we are not allowed to subtract from the concept of a phenomenal state we get through introspection. So I can see how introspection could be taken to support factualism, but only if it supports it on the strong reading.

These considerations speak for discarding WF-. But if a factualist is committed to the strong view, then she is committed to rather surprising consequences on the scientific picture of the world. For notice that the strong view directly implies *the incorrigibility of error*.

The incorrigibility of error If a phenomenal state is tokened in non-veridical conditions, it could not possibly be tokened in veridical conditions.

If (strong) factualism is correct and we are to take scientists' claims about the basic constitution of the world at face value, we should conclude not only that perceptual experience is misleading, but that our perceptual experiences could not even have

been veridical. Not only are we unlucky enough to be in a deceptive universe, but our minds are essentially deceptive.

The principal problem with the incorrigibility of error is that it flies in the face of what just about every single human being believes. All human beings take a majority of their visual experiences to be veridical. Well educated people know, in the back of their minds, that nothing is really (Edenically) colored or densely filled, but they nevertheless systematically believe their senses in everyday life. This strongly suggests that the idea that our visual experiences are veridical is fully coherent. In absence of evidence to the contrary, at least, there is every reason to think that our visual experiences could have been fully veridical. If they are also in fact massively illusory (as science shows), it follows that strong factualism is false.

There are other ways of supporting this line of argument than by appealing to the scientific image of the world. Consider blur, perspective, and gestalt effects. In chapter 6, I argued that variations in blur, perspective, and gestalt effects correspond to systematic differences in the contents of experiences. But I didn't argue that the contents of the relevant experiences are veridical in every respect, and it seems prima facie plausible that we need to posit contents which are not veridical in order to account for these phenomena. The tentative accounts of the contents of b-states and grouping experiences I offered indeed make their contents non-veridical. On the face of it, there are not enough relevant properties instantiated in the actual world to account for these kinds of perceptual variation. If blurry, perspectival, and gestalt experiences are relations to states of affairs, they are relations to non-obtaining states of affairs. This is a major problem for strong factualism given that one's total visual experience at a time always reflects the effects of some blur and some perspective: one's total experience is always necessarily non-veridical. Consider also the case of bodily experiences and emotional feelings. It is very hard to see what obtaining states of affairs these could relate us to. They seem to be necessarily non-veridical on the factualist view. The same is true of sensory

imagery. It seems undeniable that imagistic experiences are generally not relations to obtaining states of affairs, because we generate them at will independently of how the world is. They are essentially non-veridical if factualism is correct.

8.4 Summary

In chapter 7, I argued that virtualism is better supported by introspection than factualism. Virtualism is simpler and more elegant than disjunctivism, so this already makes a good case for the view. In this chapter I have discussed the two main arguments for factualism which are (mostly) independent of introspection and found them wanting. I have also argued that this view does not square well with the scientific image of the world and other considerations which show that perceptual experience is pervaded with illusion.

Nomenclature

Named propositions

Disjunctivism The phenomenal states instantiated in veridical experience cannot be instantiated in non-veridical experience. (p. 49)

Factualism The phenomenal states instantiated in veridical perceptual experience are states of standing in a factive relation to states of affairs. (p. 43)

Impure virtualism (IV) For any manner of representation M, there is a relation R such that: 1) For any basic phenomenal state *s* which has M there is some proposition P such that *s* = standing in R to P; 2) R is virtual with respect to basic phenomenal states. (p. 145)

IV- Impure virtualism is true, but pure virtualism is false. (p. 145)

Objectivism The phenomenal states instantiated in non-hallucinatory perceptual experience are states of standing in an acquaintance relation to ordinary objects. (p. 44)

Particularism The phenomenal states instantiated in non-hallucinatory perceptual experience are states of standing in a relation to states of affairs involving ordinary objects. (p. 44)

PCSE+ For every basic sensory phenomenal state *e* in modality ϕ, there is a phenomenal B-perceiving *i* in modality ϕ such that *e* = *i*. (p. 112)

Perceptual conception of sensory experience, the (PCSE) For every sensory phenomenal state e in modality ϕ, there is a phenomenal perceiving i in modality ϕ such that $e = i$. (p. 105)

Sensory virtualism (SV) For any sensory modality M, there is a relation R such that: 1) For any basic phenomenal state s in M, there is some proposition P such that s = standing in R to P; 2) R is virtual with respect to basic phenomenal states. (p. 40)

(Pure) virtualism There is a relation R such that: 1) For any basic phenomenal state s, there is a proposition P such that s = standing in R to P; 2) R is virtual with respect to basic phenomenal states. (p. 37)

(Pure) virtualism+ (Pure) virtualism is true, and relation R is the experiencing relation. (p. 143)

Weak virtualism (WV) For every basic phenomenal state s, there is a virtual relation R and a proposition P such that: 1) s = standing in R to P; 2) R is virtual with respect to basic phenomenal states. (p. 41)

Terms

B-perceiving A state which can be denoted by a predicate of the form "ϕ S", where ϕ is a perceptual verb and S is either a bare infinitive clause, a participial clause, or a small clause which complements the verb. (p. 111)

Basic phenomenal state A phenomenal state which is not derivative. (p. 36)

Cognitive experience An experience of a kind not normally associated with an emotion or a sensory process. (p. 3)

D-perceiving A state which can be denoted by a predicate of the form "ϕ S", where ϕ is a perceptual verb and S complements the verb, which is not a

B-perceiving. (p. 111)

Derivative phenomenal state A state which consists in being in one of a given set of phenomenal states distinct from itself. (p. 36)

Experience (technical sense) An event which consists in instantiating a phenomenal state. (p. 6)

Experiencing relation, the The introspectively salient relation which is at least partly constitutive of seeing$_p$, hearing$_p$, smelling$_p$, feeling$_p$, etc. (p. 142)

Intensional perceiving A state ascribable using a locution of the form "α ϕ-s S", where ϕ is a perceptual verb used intensionally and S complements ϕ. (p. 94)

Intensional reading Intensional reading A reading of an expression of the form "α ϕ-s S", where ϕ is a verb and S complements ϕ, is intensional just in case it is not material. (p. 92)

Material reading A reading of an expression of the form "α ϕ-s S", where ϕ is a verb and S complements ϕ, is material just in case its assigns the expression a logical form in which the argument of ϕ specified by S is or involves a variable bound by a quantifier located outside the arguments of ϕ. (p. 92)

Phenomenal character See *phenomenal state.* The phenomenal character of an experience is the phenomenal state it is an instantiation of. (p. 6)

Phenomenal perceiving A state ascribed by a pure phenomenal ascription. (p. 105)

Phenomenal property See *phenomenal state*.

Phenomenal state A state of the kind best exemplified by the states a) instantiated by individuals in sensory, emotional, and cognitive experiences b) individuated by the felt components they confer to such experiences. [everyday sense of "experience"] (p. 4)

Pure phenomenal ascription A perceptual ascription of the form "NP ϕ S" in which the predicate "ϕ S" denotes the state denoted by "ϕ-ly experiencing$_i$ S". (in the technical sense of "experiencing"). (p. 105)

Virtual A relation R is virtual with respect to a set S of its state instances iff there is a state *s* in S which is a state of standing in R to some *x* and which is such that it is metaphysically possible for *s* to obtain whether *x* obtains or not. (p. 35)

Subscripts

i e.g. seeing$_i$. Intensional use of a perceptual verb. (p. 94)

m e.g. seeing$_m$. Material use of a perceptual verb. (p. 94)

p e.g. seeing$_p$. Phenomenal use of a perceptual verb. (p. 99)

c e.g. seeing_c. Non-phenomenal use of a perceptual verb. (p. 103)

int e.g. modality$_{int}$, visual$_{int}$. Internal-organ modality. (p. 157)

phe e.g. modality$_{phe}$, visual$_{phe}$. Phenomenal modality. (p. 157)

rep e.g. modality$_{rep}$, visual$_{rep}$. Representational modality. (p. 157)

sti e.g. modality$_{sti}$, visual$_{sti}$. Stimulus modality. (p. 157)

Bibliography

Alston, W. P. (1999). Back to the Theory of Appearing. *Philosophical Perspectives*, 13:181–203.

Alter, T. (2006). Does Representationalism Undermine the Knowledge Argument? In Alter, T. and Walter, S., editors, *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*, pages 65–76. Oxford University Press.

Anscombe, G. E. M. (1965). The Intentionality of Sensation: A Grammatical Feature. In Butler, R. J., editor, *Analytic Philosophy*, pages 158–80. Blackwell.

Armstrong, D. M. (1968). *A Materialist Theory of the Mind*. Routledge.

Austin, J. L. (1962). *Sense and Sensibilia*. Oxford University Press.

Auvray, M. and Spence, C. (2008). The Multisensory Perception of Flavor. *Consciousness and Cognition*, 17:1016–1031.

Ayer, A. J. (1940). *The Foundations Of Empirical Knowledge*. Macmillan.

Barwise, J. (1981). Scenes and Other Situations. *Journal of Philosophy*, 78(7):369–397.

Barwise, J. and Cooper, R. (1981). Generalized Quantifiers and Natural Language. *Linguistics and Philosophy*, 4(2).

Barwise, J. and Perry, J. (1983). *Situations and Attitudes*. MIT Press.

Bayne, T. J. and Chalmers, D. J. (2003). What is the Unity of Consciousness? In Cleeremans, A., editor, *The Unity of Consciousness*. Oxford University Press.

Bealer, G. (1979). Theories of Properties, Relations, and Propositions. *Journal of Philosophy*, 76(11):634–648.

Bealer, G. and Monnich, U. (1981). Property Theories. In Gabbay, D. M. and Guenthner, F., editors, *Handbook of Philosophical Logic*, volume 10.

Biggs, S. (2009). The Scrambler: An Argument Against Representationalism. *Canadian Journal of Philosophy*, 39(2):pp. 215–236.

Blank, A. (1958). Axiomatics of Binocular Vision. *Journal of the Optical Society of America*, 48:328–334.

Block, N. (1978). Troubles with Functionalism. *Minnesota Studies in the Philosophy of Science*, 9:261–325.

Block, N. (1990). Inverted Earth. *Philosophical Perspectives*, 4:53–79.

Block, N. (1996). Mental Paint and Mental Latex. *Philosophical Issues*, 7:19–49.

Block, N. (1998). Is Experiencing Just Representing? *Philosophy and Phenomenological Research*, 58(3):663–670.

Block, N. (2003). Mental Paint. In Hahn, M. and Ramberg, B., editors, *Reflections and Replies: Essays on the Philosophy of Tyler Burge*, pages 165–200. MIT Press.

Boghossian, P. A. and Velleman, J. D. (1989). Color as a Secondary Quality. *Mind*, 98(January):81–103.

Bourget, D. (forthcoming). Consciousness is Underived Intentionality. *Noûs*.

Brentano, F. C. (1874). *Psychology From an Empirical Standpoint*. Routledge.

Brewer, B. (1999). *Perception and Reason*. Oxford University Press.

Brewer, B. (2006). Perception and Content. *European Journal of Philosophy*, 14(2):165–181.

Brewer, B. (2007). Perception and its Objects. *Philosophical Studies*, 132(1):87–97.

Brewer, B. (2008). How to Account for Illusion. In Haddock, A. and Macpherson, F., editors, *Disjunctivism: Perception, Action, Knowledge*, pages 168–180. Oxford University Press.

Broad, C. D. (1952). Some Elementary Reflexions on Sense-Perception. *Philosophy*, 27(January):3–17.

Bult, J. H. F., de Wijk, R. A., and Hummel, T. (2007). Investigations on Multimodal Sensory Integration: Texture, Taste, and Ortho- and Retronasal Olfactory Stimuli in Concert. *Neuroscience Letters*, 411:6–10.

Byrne, A. (2001). Intentionalism Defended. *Philosophical Review*, 110(2):199–240.

Byrne, A. and Hilbert, D. R. (2003). Color Realism and Color Science. *Behavioral and Brain Sciences*, 26(1):3–21.

Byrne, A. and Logue, H. (2008). Either/Or. In Haddock, A. and Macpherson, F., editors, *Disjunctivism: Perception, Action, Knowledge*, volume 29, pages 314–19. Oxford University Press.

Campbell, J. (2002). *Reference and Consciousness*. Oxford University Press.

Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.

Chalmers, D. J. (2002). Does Conceivability Entail Possibility? In *Conceivability and Possibility*. Oxford University Press.

Chalmers, D. J. (2004). The Representational Character of Experience. In Leiter, B., editor, *The Future for Philosophy*. Oxford University Press.

Chalmers, D. J. (2006). Perception and the Fall From Eden. In *Perceptual Experience*. Oxford University Press.

Child, W. (1994a). *Causality, Interpretation, and the Mind*. Oxford University Press.

Child, W. (1994b). On the Dualism of Scheme and Content. *Proceedings of the Aristotelian Society*, 94:53–71.

Chisholm, R. M. (1948). *Perceiving: A Philosophical Study*. Cornell University Press.

Chisholm, R. M. (1950). The Theory of Appearing. In Black, M., editor, *Philosophical Analysis*. Prentice Hall.

Cleve, J. V. (2002). Thomas Reid's Geometry of Visibles. *Philosophical Review*, 111(3):373–416.

Coburn, R. C. (1977). Intentionality and Perception. *Mind*, 86(January):1–18.

Crane, T. (1988). The Waterfall Illusion. *Analysis*, 48(June):142–47.

Crane, T. (2000). The Origins of Qualia. In Crane, T. and Patterson, S. A., editors, *The History of the Mind-Body Problem*. Routledge.

Crane, T. (2003). The Intentional Structure of Consciousness. In Smith, Q. and Jokic, A., editors, *Consciousness: New Philosophical Perspectives*. Oxford University Press.

Crane, T. (2005). The Problem of Perception. In Zalta, E. N., editor, *Stanford Encyclopedia of Philosophy*.

Dainton, B. (2008). The Experience of Time and Change. *Philosophy Compass*, 3(4):619–638.

Dainton, B. F. (2000). *Stream of Consciousness: Unity and Continuity in Conscious Experience*. Routledge.

Davidson, D. (1973). The Material Mind. In Suppes, P., editor, *Logic, Methodology and the Philosophy of Science*. North-Holland.

Davidson, D. (1974). Psychology as Philosophy. In Brown, S., editor, *Philosophy of Psychology*. Harper & Row.

Davidson, D. (1984). *Inquiries Into Truth And Interpretation*. Oxford University Press.

Declerck, R. (1982). The Triple Origin of Participial Perception Verb Complements. *Linguistic Analysis*, 10:1–26.

Dennett, D. C. (1971). Intentional Systems. *Journal of Philosophy*, 68(February):87–106.

Dennett, D. C. (1987). *The Intentional Stance*. MIT Press.

Dennett, D. C. (1988). Quining Qualia. In Marcel, A. J. and Bisiach, E., editors, *Consciousness in Contemporary Science*. Oxford University Press.

Dowty, D. R., Wall, R. E., and Peters, S. (1981). *Introduction to Montague Semantics*. Springer.

Dretske, F. (1981). *Knowledge and the Flow of Information*. MIT Press.

Dretske, F. (1986). Misrepresentation. In Bogdan, R., editor, *Belief: Form, Content, and Function*. Oxford University Press.

Dretske, F. (1988). Sensation and Perception. In Dancy, J., editor, *Perceptual Knowledge*. Oxford University Press.

Dretske, F. (1995). *Naturalizing the Mind*. MIT Press.

Dretske, F. (2003). The Intentionality of Perception. In Smith, B., editor, *John Searle*. Cambridge: Cambridge University Press.

Ducasse, C. J. (1942). Moore's Refutation of Idealism. In Schilpp, P. A., editor, *The Philosophy of G. E. Moore*, pages 232–3. Open Court.

Ellis, J. (2007). Content Externalism and Phenomenal Character: A New Worry About Privileged Access. *Synthese*, 159(1).

Felser, C. (1999). *Verbal Complement Clauses: A Minimalist Study of Direct Perception Constructions*. John Benjamins.

Fish, W. (2009). *Perception, Hallucination, and Illusion*. Oxford University Press.

Forbes, G. (2006). *Attitude Problems*. Oxford University Press.

Forbes, G. (2008). Intensional Transitive Verbs. In *Stanford Encyclopedia of Philosophy*.

French, R. (1987). The Geometry of Visual Space. *Noûs*, 21(June):115–133.

Godfrey-Smith, P. (1992). Indication and Adaptation. *Synthese*, 92(2):283–312.

Goldman, A. (1993). The Psychology of Folk Psychology. *Behavioral and Brain Sciences*, 16:15–28.

Gray, R. (2003). Tye's Representationalism: Feeling the Heat? *Philosophical Studies*, 115(3):245–256.

Grice, H. P. (1962). Some Remarks About the Senses. In Butler, R. J., editor, *Analytical Philosophy, First Series*. Oxford University Press.

Grice, H. P. (1988). The Causal Theory of Perception. In Dancy, J., editor, *Perceptual Knowledge*. Oxford University Press.

Harman, G. (1990). The Intrinsic Quality of Experience. *Philosophical Perspectives*, 4:31–52.

Hellie, B. (2006). Beyond Phenomenal Naivete. *Philosophers' Imprint*, 6(2):1–24.

Hellie, B. (2007). Factive Phenomenal Characters. *Philosophical Perspectives*, 21(1):259–306.

Hintikka, J. (1969). The Logic of Perception. In Hintikka, J., editor, *Models for Modalities*. Reidel.

Hinton, J. M. (1967). Visual Experiences. *Mind*, 76(April):217–227.

Hinton, J. M. (1973). *Experiences: An Inquiry Into Some Ambiguities*. Oxford: Clarendon Press.

Jackson, F. (1975). On the Adverbial Analysis of Visual Experience. *Metaphilosophy*, 6(April):127–135.

Jackson, F. (1976). The Existence of Mental Objects. *American Philosophical Quarterly*, 13(January):33–40.

Jackson, F. (1977). *Perception: A Representative Theory*. Cambridge University Press.

Jackson, F. (2003). Mind and Illusion. In O'Hear, A., editor, *Minds and Persons*, pages 421–442. Cambridge University Press.

Jackson, F. (2004). Representation and Experience. In Clapin, H., editor, *Representation in Mind*. Elsevier.

John, J. (2005). Representationism, Phenomenism, and the Intuitive View. *Philosophical topics*, 33:159–184.

Johnston, M. (2006). Better Than Mere Knowledge? The Function of Sensory Awareness. In Gendler, T. and Hawthorne, J., editors, *Perceptual Experience*. Oxford University Press.

Kim (1976). Events as Property Exemplifications. In Brand, M. and Walton, D., editors, *Action Theory*. D. Reidel.

Kim, J. (1991). Events: Their Metaphysics and Semantics. *Philosophy and Phenomenological Research*, 51(3):641–646.

Kind, A. (2007). Restrictions on Representationalism. *Philosophical Studies*, 134:405–427.

Kirk, R. E. (1996). Why Ultra-Externalism Goes Too Far. *Analysis*, 56(2):73–79.

Kriegel, U. (2003). Consciousness, Higher-Order Content, and the Individuation of Vehicles. *Synthese*, 134(3):477–504.

Langsam, H. (1997). The Theory of Appearing Defended. *Philosophical Studies*, 87(1):33–59.

Levin, D. T., Simons, D. J., Angelone, B. L., and Chabris, C. (2002). Memory for centrally attended changing objects in an incidental real-world change detection paradigm. *British Journal of Psychology*, 93:289–302.

Levine, J. (2006). Color and Color Experience: Colors as Ways of Appearing. *Dialectica*, 60(3):269–282.

Loar, B. (1990). Phenomenal States. *Philosophical Perspectives*, 4:81–108.

Lopes, D. M. M. (2000). What is It Like to See with Your Ears? The Representational Theory of Mind. *Philosophy and Phenomenological Research*, 60(2):439–453.

Luneburg, R. (1947). *Mathematical Analysis of Binocular Vision*. Princeton University Press.

Luneburg, R. (1950). The Metric of Binocular Visual Space. *Journal of the Optical Society of America*, 40:627–642.

Lycan, W. G. (1987). Phenomenal Objects: A Backhanded Defense. *Philosophical Perspectives*, 3:513–26.

Lycan, W. G. (1996a). *Consciousness and Experience*. MIT Press.

Lycan, W. G. (1996b). Layered Perceptual Representation. *Philosophical Issues*, 7:81–100.

Lycan, W. G. (1996c). Replies to Tomberlin, Tye, Stalnaker and Block. *Philosophical Issues*, 7:127–142.

Lycan, W. G. (2003). Perspectival Representation and the Knowledge Argument. In Smith, Q. and Jokic, A., editors, *Consciousness: New Philosophical Perspectives*. Oxford University Press.

Macpherson, F. (1999). Perfect Pitch and the Content of Experience. *Philosophy and Anthropology*, 3(2).

Macpherson, F. (2003). Novel Colours and the Content of Experience. *Pacific Philosophical Quarterly*, 84(1):43–66.

Macpherson, F. (2005). Colour Inversion Problems for Representationalism. *Philosophy and Phenomenological Research*, 70(1):127–152.

Macpherson, F. (2006). Property Dualism and the Merits of Solutions to the Mind-Body Problem: A Reply to Strawson. *Journal of Consciousness Studies*, 13(s 10-11):72–89.

Malone, D. R. and Hannay, J. H. (1978). Hemispheric Dominance and Normal Color Memory. *Neuropsychologia*, 16:51–59.

Martin, M. G. F. (2002). The Transparency of Experience. *Mind and Language*, 4(4):376–425.

Martin, M. G. F. (2004). The Limits of Self-Awareness. *Philosophical Studies*, 120(1-3):37–89.

Martin, M. G. F. (2006). On Being Alienated. In Gendler, T. S. and Hawthorne, J., editors, *Perceptual Experience*. Oxford University Press.

Maund, B. (2008). Color. In *Stanford Encyclopedia of Philosophy*.

Maund, J. B. (1995). *Colours: Their Nature and Representation*. Cambridge University Press.

McDowell, J. (1982). Criteria, Defeasibility, and Knowledge. *Proceedings of the British Academy*, 68:455–79.

McDowell, J. (1986). Singular Thought and the Extent of "Inner Space". In McDowell, J. and Pettit, P., editors, *Subject, Thought, and Context*. Clarendon Press.

McDowell, J. (1994). *Mind and World*. Cambridge: Harvard University Press.

McDowell, J. (1995). Knowledge and the Internal. *Philosophy and Phenomenological Research*, 55(4):877–93.

McDowell, J. (2002). Knowledge and the Internal Revisited. *Philosophy and Phenomenological Research*, 64(1):97–105.

McDowell, J. (2008). The Disjunctive Conception of Experience as Material for a Transcendental Argument. In Macpherson, F. and Haddock, A., editors, *Disjunctivism: Perception, Action, Knowledge*, volume 25, pages 19–33. Oxford University Press.

McGinn, C. (1988). Consciousness and Content. *Proceedings of the British Academy*, 74:219–39.

McGinn, C. (1989). *Mental Content*. Oxford University Press.

McGrath, M. (2008). Propositions. In *Stanford Encyclopedia of Philosophy*.

McGurk, H. and MacDonald, J. (1976). Hearing Lips and Seeing Voices. *Nature*, 264:746–748.

Millikan, R. G. (1984). *Language, Thought and Other Biological Categories*. MIT Press.

Millikan, R. G. (1989). Biosemantics. *Journal of Philosophy*, 86:281–97.

Millikan, R. G. (2004). *Varieties of Meaning*. MIT Press.

Montague, R. (1973). The Proper Treatment of Quantification in Ordinary English. In *Approaches to Natural Language*, volume 49, pages 221–242. Dordrecht.

Moore, G. E. (1905). The Nature and Reality of the Objects of Perception. *Proceedings of the Aristotelian Society*, 6:68–127.

Moore, G. E. (1910). Sense-Data. In *Some Main Problems of Philosophy*. George Allen and Unwin.

Mulligan, K. (2008). Facts. In *Stanford Encyclopedia of Philosophy*.

Nagel, T. (1974). What is It Like to Be a Bat? *Philosophical Review*, 83(October):435–50.

Neander, K. (1998). The Division of Phenomenal Labor: A Problem for Representationalist Theories of Consciousness. *Philosophical Perspectives*, 12:411–34.

Nickel, B. (2006). Against Intentionalism. *Philosophical Studies*.

Nilsson, T. H. and Nelson, T. M. (1981). Delayed Monochromatic Hue Matches Indicate Characteristics of Visual Memory. *Journal of Experimental Psychology: Human Perception and Performance*, 7:141–150.

Ono, H., Fay, A., and Tarbell, S. E. (1986). A 'Visual' Explanation of Facial Vision. 48:57–62.

O'Shaughnessy, B. (1957). An Impossible Auditory Experience. *Proceedings of the Aristotelian Society*, 57:53–82.

O'Shaughnessy, B. (1980). *The Will: A Dual Aspect Theory (2 Vols.)*. Cambridge University Press.

Pautz, A. (2006). Sensory Awareness is Not a Wide Physical Relation: An Empirical Argument Against Externalist Intentionalism. *Noûs*, 40(2):205–240.

Pautz, A. (2007). Intentionalism and Perceptual Presence. *Philosophical Perspectives*, 21(1):495–541.

Pautz, A. (2009). What Are the Contents of Experiences? *Philosophical Quarterly*, 59(236):483–507.

Peacocke, C. (1983). *Sense and Content: Experience, Thought, and Their Relations*. Oxford University Press.

Peacocke, C. (1992). Scenarios, Concepts, and Perception. In Crane, T., editor, *The Contents of Experience*. Cambridge University Press.

Perry, J. (2001). *Knowledge, Possibility, and Consciousness*. MIT Press.

Pitcher, G. (1971). *A Theory Of Perception*. Princeton University Press.

Price, H. H. (1932). *Perception*. Methuen and Co.

Putnam, H. (1975). The Meaning of 'Meaning'. *Minnesota Studies in the Philosophy of Science*, 7:131–193.

Raffman, D. (1995). On the Persistence of Phenomenology. In Metzinger, T., editor, *Conscious Experience*. Ferdinand Schoningh.

Reid, T. (2000). *An Inquiry Into the Human Mind and the Principles of Common Sense*. Edinburgh University Press.

Rey, G. (1998). A Narrow Representationalist Account of Qualitative Experience. *Philosophical Perspectives*, 12:435–58.

Robinson, H. (1994). *Perception*. Routledge.

Rosenberg, G. H. (2004). *A Place for Consciousness: Probing the Deep Structure of the Natural World*. Oxford: Oxford University Press.

Rosenthal, D. M. (2002). Explaining Consciousness. In Chalmers, D. J., editor, *Philosophy of Mind: Classical and Contemporary Readings*, volume 46, pages 109–131. Oxford University Press.

Ruys, E. and Winter, Y. (2001). Quantifier Scope in Formal Linguistics. In Gabbay, D. and Guenthner, F., editors, *Handbook of Philosophical Logic*. Harvard University Press.

Schwitzgebel, E. (2002). A Phenomenal, Dispositional Account of Belief. *Noûs*, 36(2):249–75.

Seager, W. E. and Bourget, D. (2007). Representationalism About Consciousness. In Velmans, M. and Schneider, S., editors, *The Blackwell Companion to Consciousness*. Blackwell.

Searle, J. R. (1980). Minds, Brains and Programs. *Behavioral and Brain Sciences*, 3:417–57.

Searle, J. R. (1983). *Intentionality: An Essay in the Philosophy of Mind*. Cambridge University Press.

Sellars, W. S. (1956). Empiricism and the Philosophy of Mind. *Minnesota Studies in the Philosophy of Science*, 1:253–329.

Shams, L., Kamitani, Y., and Shimojo, S. (2000). Illusions. What You See is What You Hear. *Nature*, pages 408–788.

Shoemaker, S. (1994). Phenomenal Character. *Noûs*, 28(1):21–38.

Shoemaker, S. (2000). Phenomenal Character Revisited. *Philosophy and Phenomenological Research*, 60(2):465–467.

Shoemaker, S. (2001). Introspection and Phenomenal Character. *Philosophical Topics*, 28(2):247–73.

Siegel, S. (2005a). Subject and Object in the Contents of Visual Experience. *Philosophical Review*, 115(3):355–88.

Siegel, S. (2005b). The Phenomenology of Efficacy. *Philosophical Topics*, 33(1):265–84.

Siegel, S. (2006). Which Properties Are Represented in Perception? In Gendler, T. S. and Hawthorne, J., editors, *Perceptual Experience*. Oxford University Press.

Siegel, S. (2007). How Can We Discover the Contents of Experience? *Southern Journal Of Philosophy*, pages 127–42.

Siewert, C. (1998). *The Significance of Consciousness*. Princeton University Press.

Smart, J. J. C. (1959). Sensations and Brain Processes. *Philosophical Review*, 68(April):141–56.

Smith, A. D. (2008). Translucent Experiences. *Philosophical Studies*, 140(2):197–212.

Snowdon, P. F. (1980). Perception, Vision, and Causation. *Proceedings of the Aristotelian Society*, 81:175–92.

Snowdon, P. F. (2005). The Formulation of Disjunctivism: A Response to Fish. *Proceedings of the Aristotelian Society*, 105:129–141.

Soames, S. (2005). *Reference and Description: The Case Against Two-Dimensionalism*. Princeton: Princeton University Press.

Sorenson, R. (2002). The Art of the Impossible. In Gendler, T. S. and Hawthorne, J., editors, *Conceivability and Possibility*.

Stoljar, D. (2004). The Argument From Diaphanousness. In Escurdia, M., Stainton, R. J., and Viger, C. D., editors, *Language, Mind and World: Special Issue of the Canadian Journal of Philosophy*, volume 30, pages 341–90. University of Alberta Press.

Strawson, G. (1994). *Mental Reality*. MIT Press.

Strawson, P. F. (1988). Perception and its Objects. In Dancy, J., editor, *Perceptual Knowledge*. Oxford University Press.

Sturgeon, S. (2000). *Matters of Mind: Consciousness, Reason and Nature*. Routledge.

Thesen, T., Vibell, J. F., Calvert, G. A., and Osterbauer, R. A. (2004). Neuroimaging of Multisensory Processing in Vision, Audition, Touch, and Olfaction. *Cognitive Processing*, 5.

Travis, C. S. (2004). The Silence of the Senses. *Mind*, 113(449):57–94.

Treisman, A. M. and Schmidt, H. (1982). Illusory Conjunctions in the Perception of Objects. *Cognitive Psychology*, 14:107–141.

Tye, M. (1995). A Representational Theory of Pains and Their Phenomenal Character. *Philosophical Perspectives*, 9:223–39.

Tye, M. (1996). Perceptual Experience is a Many-Layered Thing. *Philosophical Issues*, 7:117–126.

Tye, M. (1999). Phenomenal Consciousness: The Explanatory Gap as a Cognitive Illusion. *Mind*, 108(432):705–25.

Tye, M. (2000). *Consciousness, Color, and Content*. MIT Press.

Tye, M. (2002a). Representationalism and the Transparency of Experience. *Noûs*, 36(1):137–51.

Tye, M. (2002b). Visual Qualia and Visual Content Revisited. In Chalmers, D. J., editor, *Philosophy of Mind: Classical and Contemporary Readings*. Oxford University Press.

Wager, A. (1999). The Extra Qualia Problem: Synaesthesia and Representationism. *Philosophical Psychology*, 12(3):263–281.

Wagner, M. S. (2006). *The Geometries of Visual Space*. Routledge.

Westerståhl, D. (2008). Generalized Quantifiers. In *Stanford Encyclopedia of Philosophy*.

Wetzel, T. (2008). States of Affairs. In Zalta, E., editor, *Stanford Encyclopedia of Philosophy*.