A HYBRID NAIVE REALIST - REPRESENTATIONALIST THEORY OF PHENOMENAL CONSCIOUSNESS

DAVID MATHERS

CORPUS CHRISTI COLLEGE

DPHIL

PHILOSOPHY

Abstract: A Hybrid Naive Realist-Representationalist Theory of Phenomenal Consciousness

This is a Thesis about a tension in the metaphysics of phenomenal consciousness. Many philosophers have endorsed the view that phenomenally conscious perceptual experiences are transparent, in the sense that when we try to attend to those experiences, we find ourselves instead attending to the objects and properties in the external world we perceive by undergoing them. At the same time, many are attracted to the view that the items we immediately attend to on introspecting our conscious perceptions constitutively fix the phenomenal properties of those experiences. And many also believe that, in itself, whether an experience is a successful perception of an object in the world, or hallucinatory need not make any difference to the phenomenal properties of that experience. These three beliefs are in tension. For if successful perceptions get their phenomenal properties fixed by the ordinary objects and perceptible properties associated with them, there is a worry that pure hallucinations will not be able to share all and only the phenomenal properties of successful perceptions. For, arguably, hallucinations are not hooked up to the world in the right way to make ordinary physical objects available for their subjects' introspective attention. If the items which constitutively fixed the phenomenal properties of successful perceptions are therefore not available to perform this task in the purely hallucinatory case, it's not clear how pure hallucinations can come to share the phenomenal properties of successful perceptions.

In this Thesis, I first clarify exactly what the tension is here. I go on to show that it is in fact unclear whether or not the original three claims are incompatible. But I then show that there is a tension of a broadly similar sort which is definitely real. I argue that those who accept transparency should also accept an extension to the transparency thesis. This extension says that the entities associated with successful perceptions which fix their phenomenal properties are fact-like entities which have ordinary objects as parts, and which I call 'perceptible scenes'. I then argue that it follows from standard views about what phenomenal properties are that, if successful perceptions get their phenomenal properties from having such perceptible scenes associated with them, then pure hallucinations will not be able to share all the phenomenal properties of successful perceptions unless they also have their phenomenal properties fixed by associated perceptible scenes. I argue, however, that it's impossible for hallucinations to have perceptible scenes associated with them in the right way. This shows that it's impossible for a successful perception and a pure hallucination to share all and only the same phenomenal properties. In the remainder of the Thesis, I develop my own view of how the phenomenal properties of conscious perceptual experiences are fixed, which takes account of these claims about perceptible scenes. It is a hybrid view, drawing from both representationalist and naive realist view in the literature, despite representationalism and naive realism usually being thought to be rivals.

Introduction

This brief introduction is intended largely as a map of the argument of the following chapters, to be consulted to clarify the overall shape of things as the reader makes her way through the following chapters.

Recent work on phenomenal consciousness has been marked by a tension between two ideas. On the one hand, most philosophers have been committed to the view that, in some sense, whether or not an experience is a fully successful perception of an object, an illusionary misperception of an object, or completely hallucinatory, should, by itself make no difference to its phenomenal properties. This is sometimes expressed in the literature as the view that, for each successful perception there could have been a purely hallucinatory experience which shared all and only its phenomenal properties. On the other hand, it has often been argued that, at least in the case of perceptual experiences (and especially vision), our phenomenally conscious experiences are in some sense 'transparent' to introspection. What is meant by this is that when we try and pay attention to those experiences and their properties, we find that we cannot. Instead, what we find ourselves paying attention to, at least in the case of non-illusory or hallucinatory experiences, is the objects we perceive, and some of their properties. These two ideas are in tension, roughly speaking, because the way in which philosophers have typically explained what they mean by 'phenomenal consciousness' and 'phenomenal properties' talk, seems to imply that the

phenomenal properties of experiences depend on items associated with them which play what I call 'the qualia role', and the items which play the qualia role have generally been assumed to include those to which we attend when we introspect our experiences. In particular, the idea has been that two experiences differ in their phenomenal properties whenever there is a qualitative difference[1] in the qualia role playing items associated with each of them. Transparency, however, suggests that there is always a qualitative difference in the qualia-role playing items associated respectively, with a successful perception and a pure hallucination, because successful perceptions do make particular physical objects phenomenally present to their subjects, but pure hallucinations cannot. Hence transparency and the claim that qualitative differences in the qualia-role playing items associated with two experiences entail that they differ in their phenomenal properties, are together in tension with the claim that whether or not an experience is a successful perception or a pure hallucination by itself makes no difference to its phenomenal properties. For, taken together, transparency, plus the claim about qualitative differences in associated qualia role playing items entailing a difference in phenomenal properties, seem to entail that it's impossible for a pure hallucination to share all and only the phenomenal properties of a successful perception.

In this Thesis, I explore this tension, trying to first formulate it precisely, then to argue that it is real, in the sense that there is no clever way to make the three claims which make up the apparent trilemma just given consistent after all. I then tentatively propose my own solution. To put things another way, my Thesis can be seen as an attempt to work out the consequences of transparency for the metaphysics of phenomenal consciousness. What I want to see is whether transparency forces us to give up either the claim that for each successful perception there is some possible pure hallucinations which shares its phenomenal properties, or the claim that the

[1] I'll be more precise about what exactly amounts to a 'qualitative' difference in this context in the main body of the Thesis.

qualia-role playing items associated with an experience fix its phenomenal properties, and if so, which of these claims we should give up.

This Thesis has four chapters:

In the first chapter, I set up the trilemma sketched above in more detail, and show that all three of the claims which make up the trilemma, when suitably fleshed out, are, individually, widely endorsed in the literature. I then discuss three possible responses to it, other than the rejection of transparency itself, which I am bracketing here. First, holding that even given transparency there need be no qualitative difference in the qualia-role playing items associated with, respectively, a successful perception and a pure hallucination. Second denying that a qualitative difference in the qualia-role playing items associated, respectively, with a successful perception, and a pure hallucination, entails that they differ in their phenomenal properties. Third dropping the idea that for every successful perception, there is some possible pure hallucination which shares all and only its phenomenal properties. I show how recent views on phenomenal consciousness which accept transparency, can be classified as taking one or other of these three options, and that, perhaps surprisingly, the distinction between the three sorts of possible responses to my trilemma cut across the distinction between representationalism and naive realism.

In chapters two and three, I go on to explore the prospects of views which take the first horn of the dilemma, and argue that transparency is perfectly compatible with there being no qualitative difference in the qualia-role playing items, associated respectively with some fully successful perception *E* and some pure hallucination *H*.

First, in ch.2, I explore and criticize the following reasonably popular line of thought. Some philosophers hold that transparency can be reconciled with the other two parts of my trilemma in the following way. Firstly, such philosophers claim that all that is required for a pure

hallucination, *H* and a successful perception *E* to share all and only the same phenomenal properties is that those experiences have the same properties (or structured items built up from properties) as their respective associated qualia-role playing items. They then go on to argue that nothing stops pure hallucinations from making present to their subjects the sort of properties which we perceive in the world in the successful perceptual case, such as shapes and colours.

The most well known view of this sort is the type of 'representationalism'[2] advocated by Fred Dretske (1995) and, at one point (1995) Michael Tye, though Mark Johnston (2004) has also put forward a version of this view that is not representationalist. In ch.2, I criticize views which hold that properties alone, or items built up solely from properties[3], are the only sort of qualia-role playing items associated with successful perceptions, and that this allows us to hold on to all three parts of the apparent trilemma from ch.1 without inconsistency. I argue that, in fact, transparency requires that the qualia-role playing items associated with our successful perceptual experiences, include fact-like entities, called 'perceptible scenes', which have (physical) *objects* and not just *properties* as parts. I also argue that if successful perceptions have such perceptible scenes as the qualia-role playing structures associated with them, then if the structured items associated with pure hallucinations are not such perceptible scenes but rather structures constructed out of properties alone, it is impossible for pure hallucinations and successful perceptions to match in their phenomenal properties. This I argue, follows from the correct way of interpreting the requirement that two experiences match in their phenomenal properties only if there is no qualitative difference in their associated qualia-role playing items.

Ch.3 then tests the assumption that, as items which have ordinary objects as parts, perceptible scenes cannot be amongst the qualia-role playing items associated with any possible pure hallucinations. This assumption is widely shared in the literature. If it is true, then, given that ch.2 established that perceptible scenes are amongst the associated qualia role playing items in the

[2] I define this term in the next chapter.
[3] Such as Russellian propositions.

successful perceptual case, we will have established that there must be a qualitative difference in the qualia-role playing items associated with a successful perception, *E* and a (possible) pure hallucination *H*. From the latter conclusion it follows that we genuinely cannot reconcile transparency, correctly interpreted, with the other two parts of the apparent trilemma introduced in ch.1. If, on the other hand, the assumption is false, then we will be able to hold on to both transparency, and the other two parts of our trilemma, since, in principle, exactly the same perceptible scenes will be able to be the qualia-role playing items in both the successful perceptual and the purely hallucinatory case. In ch.3, I show that the view that it's possible for pure hallucinations to make ordinary objects, or items like perceptible scenes which have ordinary objects as parts phenomenally present to their subjects, is, surprisingly, far from as obviously false as it is taken to be. Eventually, however, I argue, tentatively, that the view still *is* false. I argue for this on the grounds that if we take it to be true, it's very hard to give an adequate account of what the qualia-role playing items associated with illusions or experiences which are partially hallucinatory, and partially successful perceptions, are.

Hence, together, ch.2 and 3 serve as my argument for the view that we cannot jointly hold on to all three parts of the apparent trilemma from ch.1. Since I'm bracketing the issue of whether transparency itself holds true this leaves us with two options. First, dropping the claim that for each successful perception, there is some (possible) pure hallucination which shares all and only its phenomenal properties. Second, dropping the claim that two experiences share all and only the same phenomenal properties iff there is no qualitative difference in the qualia-role playing items associated with each of them. In ch.4, I address the question of which option to take here, and defend the view that we should drop the first, rather than the second of these claims. I then go on to develop my own view which is consistent with the results of ch.s.2 and 3. My view effectively identifies phenomenal properties with properties of consciously successfully representing objects to have properties (or to be related to each other in certain ways). I hold that the qualia-role playing structure associated with a successful perception is the perceptible

scene determined by the representational content of that successful perception. In the rest of ch.4 I first outline some motivations for accepting my view over some similar rival views. I then show that, despite its resemblance to both some forms of representationalism, and some forms of naive realism, my view is not vulnerable to a range of either standard naive realist arguments against representationalism, or standard representationalist arguments against naive realism. I end ch.4 by finally discussing one possible problem for my view if illusion turns out to be extremely pervasive.

Before we get started though, I need to give a brief summary of how I am using key terms.

Firstly, 'successful perception', 'illusion' and 'hallucination'. I will assume that we perceive both objects and their properties. I understand successful perception as perceiving both an object and one or more its properties (or how two or more objects are related). Illusion occurs when we perceive an object, but misperceive one or more of properties. Hallucinations are cases where humans or other perceiving subjects undergo an experience as of perceiving an object without thereby actually perceiving any object. I take it that the distinction between perception, illusion, and hallucination is embedded in ordinary thought about perception. While it is certainly possible to question its cogency, I won't do so here. Of course, there can be an experience that involves accurate perception of the properties of one object, misperception of another (or another property of the same object) and hallucination of a third. So a *single* experience can theoretically be a successful perception of one object, *o* and some of its properties, whilst also being an illusory misperception of either *o* or some second object. Despite this, when I speak of 'successful perceptions' as a class, I will, unless I say otherwise, mean the class of fully successful perceptions, experiences which are both successful perceptions of objects, and also not illusory misperception of any objects, or hallucinations of any objects. I will however, at times, speak of particular experiences being both fully successful perceptions of some object and one or more of

its properties, and also illusory or hallucinatory. But references to successful perceptions as a class, are to be read as being to *fully* successful perceptions, unless I indicate otherwise.

Within the class of hallucinations, we can distinguish the *impure* hallucinations from the *pure* hallucinations. The former are hallucinations which are also perceptions or misperceptions of some object, that is hallucinations which also involve successful perception and illusion. The latter are hallucinations that are just hallucinations, and not do not involve the perception or misperception of any particular objects. So for example, if I hallucinate a wardrobe whilst perceiving a spider on top of it, this would be an impure hallucination, whilst if I hallucinate both the spider and the wardrobe and they fill the whole of my visual field, this would be an pure hallucination. There may never have been a pure hallucination, but I will rely here only on the assumption that they are nomologically possible.

Note that 'perceptual success', in the sense which I'm using the term is a different thing from 'veridicality'. Veridicality is generally taken to be a matter of things being as they appear. But things can be as they appear, through a massive coincidence, even if I fail to perceive anything, and an object can appear to me to have a property which it actually has, even though I do not successfully perceive it to have that property. For example, I might hallucinate a tomato, and by coincidence, there might be a tomato before me which looks exactly like the one I hallucinate. Similarly, the following case is obviously at least metaphysically and probably nomologically possible. I see a tomato, *t*, and it appears red to me. But whilst my seeing the tomato, and its roundness, is caused by the tomato reflecting light into my retina, my seeing it as the shade of red it actually is, is in fact caused by the interference of a tiny nanobot within my brain. Success in perceiving both objects and their properties depends on an experience being caused in the right way, but veridicality does not[4].

[4] I have drawn heavily here on Siegel's (2010, pp.34-9) discussion of these issues.

Secondly, throughout this Thesis, I will need to speak of the things which play a constitutive role in fixing the phenomenal properties of our experiences, without begging the question as to what ontological categories these things fall into. In particular, I will want to refer to the items associated with our experiences which play a constitutive role in fixing their phenomenal properties, without implying that they are objects, or properties, or tropes. I will use 'item' as a very general term which covers things of any ontological category whatsoever: particular objects, such as tables and chairs, or for that matter mental particulars like sense-data if such there be, properties, like *being red* or *being circular and five feet from the Eiffel Tower*, and tropes of properties. Hence, on my views, tables, angels, sense-data, and numbers, are all putative objects, whilst properties and tropes of properties are not 'objects' in my sense.

Thirdly, and finally, although I will take the claims argued for in this Thesis to apply to phenomenally conscious perceptual experiences as a whole, I will generally focus on the case of visual experiences, specifically. At times, for ease of exposition, I'll even speak as though all perceptual experiences have features which, strictly speaking belong to visual experiences alone. For example I'll speak of our perceptual experiences as making colours 'phenomenally present' (this term gets defined in ch.1), to their subjects, when it is really only visual experiences that do this. So if a claim obviously only applies to visual experiences 'perceptual experience' should be read with a tacit restriction to visual experiences. I am aware that there are dangers in taking vision as a model for all sense modalities, but at the same time, the vast majority of the literature I am engaging with here, mostly centres on vision, with claims then transferred to the other senses. Furthermore, most of this Thesis will be critical, and if views of perceptual experience in general are inadequate in the visual case, then they are unacceptable, regardless of whether the arguments that show they are inadequate could in principle be re-worked to apply in, say, the tactile case. When I do tentatively put forward a view of my own in ch.4, I put it forward as a proposal about the metaphysics of phenomenally conscious perceptual states, not just *visual* phenomenally conscious perceptual states. But I am happy to admit that the evidence I give for

the view comes largely from reflection on vision alone, and that it's possible there are problems for the view which only emerge when we look at what's going on in the case of other sorts of perceptual experiences. At least if the view does better on the visual case than its competitors in the literature, it has something going for it.

Chapter One: Transparency and Hallucinatory Matching; Is there a Tension?

Section One-Introducing the Guiding Question

In this thesis, I want to tackle a tension that emerges between two popular ideas in the philosophy of perception, when a third idea implicit in much recent work is brought out into the open. The two popular ideas are the following. Firstly that for each successful perception there is a possible (pure) hallucination which is in some sense an experience of the same type, and secondly that perceptual experience is 'transparent' in that when we introspect our experience we find we are not directly aware of the experiences themselves or their properties, but rather, at least in the case of successful perception, of the external objects the experience is an experience of, or perceptible properties of such objects. These two ideas, taken together, are in tension with a third idea implicit in much of the recent literature, which is that the type of experience a particular perceptual experience is, depends on the items that that experience makes available for introspective attention[5]. A somewhat similar tension has already been noted by Michael Martin, particularly in (forthcoming, ch.1 and ch.3), although the details of how I

[5] This is a very rough formulation, which I'll make much more precise below.

formulate each of the three ideas differs. In the rest of this Thesis, I discuss how to resolve the tension. But this opening chapter, I just want to explain the three ideas, the tension between them, and the different possible options for resolving this tension.

We can start by precisely formulating the idea that for each successful perception there is some metaphysically possible hallucination which is an experience of the same kind, what I refer to as the 'Hallucinatory Matching Thesis'. As I see it, this should be formulated in terms of there being, for each successful perception, a possible[6] hallucination which is the same kind of phenomenally conscious state. The reason for this is that, amongst the ways in which hallucinations have generally been taken to be able to match successful perceptions, it is mostly clear sameness in phenomenal properties which the literature has taken to depend on a match between the items which two experiences make available for introspective attention. It is therefore only when formulated in terms of phenomenal properties that the Hallucinatory Matching thesis is obviously in tension with the claim that it is only ordinary physical objects and/or perceptible properties of ordinary objects which our perceptual experiences make available for introspective attention. When understood in terms of a claim about phenomenal consciousness, the Hallucinatory Matching Thesis can be formulated as follows:

Hallucinatory Matching Thesis (HMT): For any successful perceptual experience *E* there is some metaphysically possible purely hallucinatory experience, *H*, which shares all and only the phenomenal properties of *E*.

Note that in formulating the HMT in terms of pure hallucinations alone, I am not meaning to imply that there is no tension between transparency, and analogous claims about impure hallucinations or illusions. Ultimately to vindicate our background intuition that whether or not an experience is a fully successful perception by itself need make no difference to its phenomenal

[6] From now on, reference to necessity and possibility should be read as implying metaphysical possibility and necessity, unless otherwise indicated.

properties, we would need to vindicate analogous matching theses for impure hallucinations and illusions as well. But I shall concentrate on the HMT here. If the HMT is false then clearly the background 'no difference' intuition is misleading at least in the important case of pure hallucinations. When I put forward my own tentative positive view of the metaphysics of phenomenal consciousness in ch.4, however, which entails the falsity of the HMT, I will explain how this view accounts for illusions and impure hallucinations as well as pure hallucinations and successful perceptions, and draw out its implications for the analogous matching theses for illusions and impure hallucinations.

Note that the way I have formulated the HMT above may or may not be the formulation that best expresses the basic intuition that the phenomenal properties of a phenomenally conscious perceptual experience don't depend on whether or not it is a successful perception. In particular, the claim that there is, for each successful perception, some metaphysically possible pure hallucination, which matches it phenomenally, is quite weak. It may be that views which accept this, but which deny that, for each successful perception, there is a matching pure hallucinationwhich meets some further restricting condition, will also strike us as counterintuitive. If so, then to really capture our convictions here, we would need a claim stronger than the HMT itself which said more than that for each successful perception, there was a matching pure hallucination. It might instead say that for every successful perception there was a matching pure hallucination which met some further condition, or even that for each successful perception, all pure hallucinations which shared some further property would match it. For example we can make the HMT into a stronger claim by making the modality involved nomological rather than metaphysical, and it certainly seems odd to deny that for each successful perception, there is a matching *nomologically* possible hallucination. We do not think that in order for there to be a pure hallucination which matches a successful perception, the laws of nature would have to be different. And there are even stronger claims, whose denial might well strike us as counter-intuitive. For example, for each successful perception, *E*, there is the counter-

factual claim that had it been somehow turned into a pure hallucination, by some powerful technology or magic breaking the link between experience and world, whilst keeping the neural base of the experience exactly intact, *E*'s phenomenal properties would not have changed. I suspect that many people who find the truth of the HMT 'obvious' will also find the claim that such counter-factuals are usually true, to be obvious. Perhaps it's commitment to the latter claim which underlies our commitment to the HMT as formulated above. An even stronger claim which might underlie our commitment to the HMT is the claim that necessarily: if two experiences, E_1 and E_2 share all and only the same narrow properties, then they share all and only the same phenomenal properties[7]. This claim I will refer to as 'phenomenal internalism'. So there is a worry that perhaps the HMT, as formulated above does not in fact best capture the relevant intuition here.

However, while it may well be that denying one or more of these stronger claims registers as just as counter-intuitive as denying the HMT (in its official formulation given above) itself, or even that our commitment to the HMT in some sense depends on our commitment to one of these claims, or another stronger claim that entails it, in practice the HMT itself is still of interest for the following reason. Any plausible way of working out the claim that whether or not an experience is a successful perception or a pure hallucination need make no difference to its phenomenal properties, will surely *entail* the HMT. Further, my main purpose in this Thesis is to argue against this claim, in ch.s2 and 3, and then develop a view which is inconsistent with it, but has other virtues, in ch.4. It therefore makes sense for me to treat the weakest reasonable formulation of the claim as the canonical one, since if I can give reasons to reject it I will thereby also give reasons to reject all the stronger formulations which entail it.

[7] Though this claim is explicitly denied by some who endorse the HMT itself, such as Michael Tye. See his (2009, ch.8), for denial of phenomenal internalism and his (2002, p.140) for the affirmation of the HMT (or something very like it).

Against the traditional view, embodied in the HMT, we can set the currently popular claim that conscious experiences are, in a sense 'transparent'. On this view, when we try to attend to the phenomenology of our conscious experiences[8], what we attend to is simply what we perceive by undergoing that experience, or at least some subset of the items we perceive. On such views, we come to know the phenomenal properties of our experiences, by attending to the things they are experiences of. The tension between this view and the HMT arises for the following reason. Part of the traditional picture around phenomenal consciousness is that there are items which play what I call the 'qualia role'. And part of what it is to play the qualia role is be an item associated with an experience which is both accessible to introspection, and yet also plays a role in constitutively fixing the phenomenal properties of that experience. As traditionally understood, 'qualia' were supposed to be narrow properties of conscious experiences themselves. But if we accept transparency, as many recent writers on phenomenal consciousness do, then no such properties of experience are available for introspective attention. Many recent writers who do accept transparency have instead suggested that something like the qualia role, and in particular the role of constitutively fixing the phenomenal properties of experiences, is played by the items that our perceptual experiences are perceptions of, rather than by properties of the experiences themselves. At this point, however, the possibility of a tension with the HMT arises. Suppose that a successful perceptual experience has the phenomenal properties it does because of the items it makes available for its subject's introspective attention, and that those items are worldly items, such as ordinary physical objects and their properties. And suppose also that purely hallucinatory experiences cannot make any such items available to the introspective attention of their subjects. It then becomes natural to worry that pure hallucinations cannot match successful perceptions phenomenally. After all, the constitutive grounds of the possession of phenomenal properties by any successful perception will, by hypothesis, not obtain in any purely hallucinatory case.

[8] Some weaker versions of the view are restricted to perceptual experiences.

In the rest of this chapter I set the stage for the following chapters by first fleshing out this case for a tension between the HMT and transparency, and then taxonomizing the possible lines of response to it. I won't here be concerned with arguing for the transparency thesis or the HMT. In the case of the HMT, I think that there is not much further evidence that can be given, beyond the fact that that its denial is highly counter-intuitive. As for transparency, the case for it has been made so often by other writers that I don't think much would be gained from a further defence of it here. When I come to discuss transparency in detail in Section 3, I will quote from three of the most famous defences of transparency, but I shan't try and give a further argument for it.

In the next two sections, I firstly give an account of how phenomenal vocabulary in its present usage was originally introduced into modern philosophy of mind, and then an account of the recent history of the transparency thesis. With this background in place, in Section Four, I give a much fuller and more precise account of why one might think that the HMT and transparency conflict, and in doing so also flesh out the currently somewhat more shadowy claim that the phenomenal properties of experiences are fixed by, roughly speaking, the items experiences make available for introspective attention. I show that these three claims together are jointly in tension. Finally, in Section Five, I give an account of the various ways in which someone might respond to this tension, and show where various current accounts of phenomenal consciousness fall within this taxonomy. This sets the stage for the rest of my thesis, where I discuss whether any of the ways of reconciling the HMT and transparency discussed in Section Five are less costly than dropping the HMT itself.

Section Two-A Brief History of 'Phenomenal Consciousness'

In this section, I'll explain some strands in talk about phenomenal consciousness which I think act as an implicit definition of the term, at least in conjunction with the lists of examples of phenomenally conscious states which also tend to be given when philosophers are explaining what they mean by 'phenomenal consciousness' and 'phenomenal properties'. In doing so I'll also say a little more about so-called 'qualia', the items that, in the classic papers and books on phenomenal consciousness, were supposed to be entities which played the dual role of being both available for introspective attention, and somehow constitutively fixing the phenomenal properties of the experiences which they were associated with.

Although the term 'phenomenal consciousness' only began to be routinely used in its current sense in the 90s, the origins of the modern use of the term, lie largely in a particular family of objections to physicalism, which arose in the 70s and early 80s[9]. These objections centred around the idea of there being some sort of 'explanatory gap'[10] between the physical on the one hand, and conscious experiences on the other. This gap was then claimed to be a problem for physicalism. For example, Kripke (1980) argued that mind-brain identity theory struggled to deal with the fact that identity statements are necessarily true if true at all, and yet it seems that, even if the conscious mental state *being in pain* is invariably correlated with a particular neurological state, there *could* have been pain even without the existence of the neural state in question, or vice versa. Nagel (1974) meanwhile, worried that it appears that one can know all the physical and neuroscientific facts about a creature, and yet not know all there is to know about its conscious life, in particular not know what it's like to be that creature[11]. A similar worry was turned into an explicit argument against physicalism by Jackson (1982). According to Jackson, someone who had never seen a coloured object could learn all the physical facts about human colour perception. But still, there would be something about perceptual experiences of, say, red

[9] At least, the modern discussion of such objections begins then. Objections of the same general type go back much further.

[10] The phrase was coined by Levine (1983), though Levine is there cautious about whether the gap really does raise problems for physicalism.

[11] Though in (1974) Nagel explicitly avoids committing to the claim that physicalism is therefore false.

which such an individual could not know without undergoing them, namely, 'what it's like to see red'. Jackson concluded that there are non-physical facts about conscious experiences which can only be known by undergoing those experiences. And Block and Fodor (1972) raised a much older worry, namely that functionalist theories run into trouble with the Lockean hypothesis of the inverted spectrum, because it's not clear why one of two functionally and computational identical minds could not be spectrum inverted with respect to another[12].

Although Jackson, and Block and Fodor make use of the related term 'qualia' these works do not use the term 'phenomenal consciousness'. Instead, they typically give examples of particular types of conscious state which they claim raise epistemic-gap related troubles for physicalism, whilst implicitly or explicitly implying that these trouble arise for the class of conscious states in general. Later, however, Chalmers (1996, ch.1) and Block (1995) (amongst others) made the claim that the natural language term consciousness is not univocal, but instead is ambiguous between two or more distinct or related readings, and it was usually claimed that it was only on one reading of 'consciousness' that it denoted the phenomena involved in the problem-cases for physicalism mentioned in the previous paragraphs. They introduce the term 'phenomenal consciousness' to refer to consciousness in this sense. In these later works, such as those by Block and Chalmers, the notion of phenomenal consciousness is usually introduced partly by examples of phenomenally conscious states (usually perceptual or affective states, but also sometimes occurrent thoughts, imaginative states, etc.), and partly by reference to the (supposed) troubles of physicalism mentioned in the previous paragraph. In other words phenomenal consciousness is said to be the sort of consciousness which appears hard to explain in physical terms, whose absence in a world where all the physical and functional facts are as actual appears to be conceivable, about which we have an intuition that Jackson's Mary would learn something when she leaves her room, and so on.

[12] Note that Block and Fodor do not reject 'physicalism' in the current sense of the term in (1972), merely functionalism.

It's also in the 90s that people started to intensively discuss the question of whether the phenomenal properties of mental states supervene on their physical properties. As well as the most general phenomenally conscious mental state *being phenomenally conscious*, discussion of this issue involved talk of specific sorts of phenomenally conscious experiences such as *being a visual experience of red* or *being in pain*. From this, talk not just of a general phenomenon 'phenomenal consciousness', but of phenomenal properties which divide states which bear the generic property *phenomenally conscious* into specific types, became common in the literature.

Having discussed 'phenomenal consciousness' (and 'phenomenal properties') let's move on to 'qualia'. Now, the term 'qualia' is widely used in the literature which I have just surveyed. Its use, however, is somewhat complex, and it's hard to give a quick account of it. Basically, qualia were supposed to be items, usually properties, associated with phenomenally conscious experiences, which were somehow tied in an important way to the explanatory gap, and to the phenomenal properties of those experiences Some philosophers *seem* to treat 'qualia' and 'phenomenal properties' as synonyms, though it's hard to find the claim that they are made completely explicitly[13]. In other cases, it's less clear that the two terms are to be regarded as synonymous, but the qualia associated with an experience are still taken to fix its phenomenal properties. So, for example, consider Levine's use of the term 'qualia' in his important (1980) article which introduced the term 'explanatory gap' into the literature. Levine endorses Kripke's claim that for any brain state claimed to be identical with the mental state *being in pain* it will be conceivable that some subject is in that brain state without being in pain. And what Levine takes himself to be endorsing is the conceivability of the absence of pain *qualia* in the presence of the appropriate brain state. Similarly Block and Fodor (1972) treat cases of spectrum inversion as cases where the qualia associated with colour experiences are switched between different subjects. Talk of

[13] Although Chalmers (1996, p.4) and Block (1996) come close.

particular sounds, scents, smells, tastes, and arguably looks[14], is supposed to be talk about particular qualia[15]. In other words, qualia-believers typically assumed that, for each determinate perceptible property there was a unique type of qualia, associated with all and only successful perceptual experiences of that property. Furthermore, it was assumed that you can know which qualia are associated with your own experiences just by introspection. And finally, qualia were supposed in some sense to fix the phenomenal properties of the experiences they were associated with[16]. Here, the fixing is supposed to involve some sort of constitutive connection, or perhaps even the identity relation holding between qualia and phenomenal properties, rather than mere causation.

Section Three-The Rise of Transparency

One feature of the above story about phenomenal consciousness and qualia that has been challenged repeatedly in the literature, is the idea that there are items which play all of the following parts of the qualia-role: they are associated with experiences and fix their phenomenal properties, we can attend to them when we pay attention to our experiences, and they are properties of experiences. For, it has been frequently claimed, when we introspect our experiences, and try to determine what those experiences are like, try to determine their phenomenal properties, we do not attend to experiences themselves or their properties or any

[14] The tradition definitely assumes that there are visual qualia, paired with particular perceptible properties, just as much as auditory, olfactory or gustatory qualia, but the analysis of 'looks' talk is a controversial issue.

[15] On this, see, once again Chalmers (1996, ch.1) and even more explicitly Dennett (1988). Dennett believes qualia talk to be fatally confused, but I think that believers in qualia would accept the particular claim just cited.

[16] This can be seen for example, in the dual facts that a) inverted spectrum scenarios were routinely described in the literature as cases where perceptions of objects with the same light-reflectance profiles lead to experiences with different qualia in different subjects (or in other words, as cases of 'qualia inversion', as in Block and Fodor (1972) and innumerable others) and b) it was taken as obvious that inverted spectrum scenarios were cases in which the phenomenal properties of experiences varied along with the inversion.

other mental items. Instead our attention seems to be drawn to items of the sort which we can perceive in the external world. In this sense, our experiences are 'transparent'; when we try to focus on them, we end up 'looking through' them at the world. Advocates of the transparency thesis[17] all agree that, in the case of successful perceptions, when we introspect them, our attention goes to the sort of items we can perceive in the world: ordinary objects, and/or their perceptible properties (or tropes thereof), such as colour and shape properties. They disagree about exactly what, if anything, we would attend to in the illusory and the hallucinatory case. But they all agree with the *negative* claim that when we introspect illusions or hallucinations, and try to attend to the items which play the role 'qualia' in the traditional sense were supposed to play in fixing their phenomenal properties, there is an important sense in which we do not attend to any mental items.

The lesson that many advocates of transparency draw from it is that the role that 'qualia', traditionally conceived of as mental items, were supposed to play in constitutively fixing the phenomenal properties of our experiences, is actually played by the perceptible items which we attend to when we introspect our experiences, and try to determine their phenomenal properties. Or at least they have held that those items should take over this role in the successful perceptual case; advocates of transparency disagree about whether pure hallucinations can bear phenomenal properties and if so, how they come to bear them. That is, they have held that, for successful perceptions, their phenomenal properties are constitutively fixed by the ordinary objects and/or perceptible properties that we attend to when we introspect them, the very items which we perceive in the world by undergoing the experiences in question. This can be seen from the fact that many advocates of transparency adopt some sort of relational account of phenomenal consciousness or at least of the successful perceptual case. On these relational accounts the phenomenal properties of experiences are fixed by exactly which ordinary objects or

[17] Advocates of transparency include, non-exhaustively, Harman (1990), Dretske (1995), Tye (1995, 2000), and Martin (2002a).

properties of ordinary objects are on the other end of some relation, *R*, which the experience bears to ordinary objects and/or their properties. Some have adopted representationalism, the view on which the phenomenal properties of experiences supervene on their representational content. On such views, the items we attend to when we introspect our perceptual experiences are simply the properties that they represent as instantiated in the world. Michael Tye (1995, 2002), is one prominent defender of this sort of view. On the other hand, some transparency lovers, have adopted so-called 'naive realist' views of phenomenal consciousness. On these views to have a phenomenally conscious experience is to have an experience which bears a special relation, *R*, to ordinary physical objects, where the phenomenal properties of experiences are determined by which objects they bear *R* to and which properties of those objects the experiences are somehow picking up on. *R* is supposed to be a relation which experiences can bear only to objects they count as perceptions of.

Here I want to get clear about how best to formulate transparency itself. Standard accounts of transparency often formulate it as a claim about all of the items which are objects of attention when we introspect, or about what we 'directly' attend to, but I shall argue that such accounts are actually problematic, and suggest an alternative way of capturing the basic idea that when we turn our attention to the items which fix the phenomenal properties of our experiences, we find that those items are ordinary objects and/or their perceptible properties, rather than mental items.

Here then, is some of what the advocates of the transparency thesis have to say for it:

> When you see a tree, you do not experience any features as intrinsic features of your experience. 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, p.39).

> If you are attending to how things look to you, as opposed to how they are independently of how they look, you are bringing to bear your faculty of introspection. But in so doing, you are not aware of any inner object or thing or event. When you introspect your visual experience, the only particulars of which you are aware are the external ones making up the scene before your eyes.

You are not aware of those objects and a further inner object or episode. ...Your experience is thus transparent to you. When you try to focus upon it, you 'see' right through it...to the things apparently outside and their apparent qualities. But when you introspect, you are certainly aware of the phenomenal character of your visual experience. ...Via introspection, you are directly aware of a range of qualities which you experience as being qualities of surfaces at varying distances away and orientations and thereby you are aware of the phenomenal character of your experience. (Tye, 2002, p.139).

When I stare at the...lavender bush at the end of my street, I can attend to the...colours and shapes of leaves and branches... . But I can also reflect on what it is like for me now to be staring at the bush, and in doing so I can reflect on particular aspects of the visual situation: for example at this distance of fifty metres the bush appears more flattened than the rose bush which forms the boundary of my house with the street. When my attention is directed out at the world, the lavender bush and its features occupy centre stage. ...when my attention is turned inwards instead to my experience, the bush is not replaced by some other entity belonging to the inner realm of the mind... . I attend to what it is like for me to inspect the lavender bush through perceptually attending to the bush itself while at the same time reflecting on what I am doing. So it does not seem to me as if there is any object apart from the bush for me to be attending to or reflecting on while doing this. (Martin, 2002a, pp.380-1).

Now the most obvious way to formulate the claim being defended by at least Harman and Tye in these quotes is that it is literally impossible, whilst introspecting them, to attend to our successful perceptual experiences and their properties. However, if we formulate transparency in this way, it is open to the following objection. If I come to know on the basis of introspection that I am having an experience of a particular kind, and I do so by focusing on what sort of experience I'm currently having, am I not in some sense attending to my experience and its properties? Certainly, there are reasonable accounts of attention on which I will, in such a case, count as attending to those items. There seems to be one reasonable ordinary sense of 'attends to' in which I attend to something if I think about it, and concentrate hard on doing it. Furthermore, at least some current philosophical accounts of attention support the idea that you are attending to something just by concentrating thought on it[18]. So the worry is that if transparency is defined as the view that we never attend to our phenomenally conscious perceptual experiences, or their properties, whilst introspecting those experiences, then transparency might well be inconsistent with the

[18] See Mole (2011) for one such account.

undeniable fact that when we introspect such experiences, we devote cognitive resources to thinking about them[19].

The last of our three quotes, from Martin, suggests a more reasonable interpretation of transparency. This is the claim that, although we do attend to our conscious perceptual experiences and their properties (including their phenomenal properties), whilst introspecting those experiences and trying to determine what phenomenal properties they have, we somehow do so *via* paying attention to the items in the external world which those experiences count as perceptions of.[20]. In other words, transparency would become the claim that the only items which we can in some sense 'directly' attend to when we introspect our successful perceptual experiences, are the ordinary physical objects or properties in the external world which they count as perceptions of. On such a view we do attend to our perceptual experiences and their phenomenal properties, when we introspect those experiences, and thereby form beliefs about

[19] Tye (2014) suggests that transparency could be formulated as the claim that it is only perceptible items, and not experiences and their properties of which we have *de re* as supposed to *de dicto awareness* when we introspect. (I don't think his formulating things here in terms of awareness rather than attention makes much difference) But this doesn't seem to help, since, presumably, I can come to think singular thoughts about my experience, thoughts which could be expressed using demonstratives like 'this experience'. Why doesn't that count as de re attention to my experience? Tye also alludes (2014, p.42) to Dretske's (1999) distinction between 'fact awareness' and object 'awareness', and claims that we are 'fact aware' but not 'object aware' of our experiences. Now, I find the distinction Dretske's makes between 'fact' and 'object' awareness pretty murky. I won't attempt to give a full argument for that here for space reasons. But briefly: Dretske doesn't seem to think of object awareness in (1999) as involving some kind of primitive 'knowledge by acquaintance' which isn't a matter of representing an object. Indeed, given that he thinks that 'all mental facts are representational facts' (1995, p.xiii) Dretske cannot mean anything like this by object awareness. But if both being object aware of an experience, and being fact aware of a fact about it, are just matters of representing the world in some way, it's hard to see what the difference between 'fact' and 'object' awareness amounts to. If 'all mental facts are representational facts' then fact awareness must surely be representing an object to have a property, and this representation by accurate (and maybe also veridical in the perceptual case.) But then object awareness, if it's a matter of representation, looks like it also has to be a matter of representing an object to have a property (unless there are representational mental states which represent objects to simply have bare properties which seems doubtful. This seems to either (if we treat object-awareness as factive) collapse object awareness into fact awareness, or (if we don't treat object awareness as requiring accuracy in the properties an object is represented to have) collapse the distinction between object and fact awareness into the distinction between accurate and inaccurate representation of an object. Either way, there will be no sense in which we can be 'fact aware' of our perceptual experiences through introspecting them, but cannot be object aware.

[20] Tye's talk of us attending to the phenomenal *character* of our experiences by attending to the things they are perceptions of my also sound like it suggests the claim that we attend to the phenomenal *properties* of our successful by attending to the items they are perceptions of. But in fact this is mistaken, since for Tye, 'phenomenal character' is not a term for the phenomenal properties of a perceptual experience, but rather for the qualia-role playing items who's association with the experience fixes their phenomenal properties.

them. But this attention is in some sense derivative of our attention to the items the experience counts as a perception of.

Now, *maybe* there is a way to formulate transparency in terms of a notion of a distinction between direct and indirect attention, which produces a true claim. However, this is far from obvious, since the notion of attending to something via attending to something else strikes me as a somewhat murky one. It's not like 'direct attention' is a term which figures in enough philosophical theorizing, independent of discussions of the transparency thesis for us to have an independent grip on it. Maybe 'direct *perception*' does, but perception and attention are not the same thing. Perhaps, however, 'direct attention' can be cashed out as simply involving some kind of causal or evidential dependence of the beliefs we form about the phenomenal properties of our experiences through introspection, on how our experiences represent the outside world to be[21].

Even then, however, there are problems. Firstly, perceptual attention is not 'direct' either, if to 'directly' attend to something is to attend to it via undergoing a representational mental state that hasn't been arrived at inferentially, at least if we count transition from unconscious states in early vision, to successful perceptions, as inferential. Still, this can probably be finessed by relativizing, so that the claim is not that we directly attend to perceptible items when we introspect, and only indirectly attend to experiences and their property, but rather that we attend to our successful experiences, and their properties, via attending to some of the properties and/or objects in the world which they count as perceptions of. Still, the notion of attending to something via attending to something else, seems just as murky and in need of further clarification as the notion of direct attention, though once again, we might be able to simply gloss

[21] Something like this way of cashing things out is suggested by the famous 'displaced perception model of introspection' developed by Dretske (1995, pp.41-4) which is sometimes associated with the transparency thesis.

it in terms of inference from perceptual representation of the world, to occurrent thoughts about perceptual experiences.

More seriously there is a further problem with framing transparency as the claim that when we introspect, we can only attend to the properties of our successful perceptual experiences, *via* attending to the items they are perceptions of. The negative part of the transparency thesis, the sense in which we do not attend to the items we perceive, is supposed to hold in the case of hallucinations and illusions as well.[22] In particular, some naïve realists, such as Martin (2004), deny that pure hallucinations[23] have any 'substantive' mental properties at all, holding instead that their only interesting mental properties are properties of being indiscriminable from successful perceptions by introspection alone. Presumably such naïve realists are going to deny that we attend to any perceptible items in the world at all in the purely hallucinatory case. After all, if such experiences lack representational content, and are not perceptions of anything, it's not clear how it could be determined which objects or perceptible properties they made available for introspective attention. If the only reason, in the successful perceptual case, why coming to form judgment about the properties of our experiences through introspection, doesn't count as 'direct' attention to our experiences, or their properties, is that such attention is somehow grounded in attention to the things we perceive, then it looks like forming introspective judgments about pure hallucinations *will* count as directly attending to them on Martin's view of the hallucinatory case. But a formulation of the transparency thesis which counts Martin as denying transparency of pure hallucinations, on the grounds that we do attend to them directly when we make

[22] For example Tye (2002), says, after giving the defence of transparency quoted above that 'the points made so far do not require your experience to be veridical…[i]ndeed the case could be one of complete hallucination' (p.140). Technically Tye's commitment here is to both the positive and the negative parts of the transparency thesis holding when we hallucinate, but this of course entails the weaker claim that the negative part holds.

[23] Actually Martin technically only commits, in (2004), to this claim for a certain subset of pure hallucinations, but since the negative part of the transparency thesis is supposed to hold for all pure hallucinations, this makes no difference.

introspective judgments about them, looks like a bad formulation, given that Martin is one of the most prominent advocates of transparency.

In the face of these difficulties, I think we should abandon attempts to define transparency in terms of some restricted sense of attention on which we can be said to attend to the items we perceive, when we introspect our successful perceptual experiences, but also such that we can never be said to attend in that sense to experiences themselves and their properties, when we introspect perceptual experiences of any sort[24]. Instead of defining transparency in terms of a particular type of *attention*, I think we should look at capturing the sense in which our experiences are 'transparent' to introspection, in terms of what we find and don't find, when we introspect are experiences, *and attend to all of the items which meet a certain criterion.* In particular, my idea is this. Most discussions of transparency, including those from Harman and Tye quoted above, occur in a context where it is being argued that there are no 'qualia' in the traditional sense, where the readers are being expected to already have some grasp of qualia-talk and phenomenal property talk, picked up from the literature described in the previous Section. So most discussion of transparency occur in a context where what we find when we introspect is supposed to be relevant to the question of whether there are qualia, conceived of as properties of experience or other mental items. Perhaps therefore, when people reading passages making an intuitive case for transparency, like those above, follow the instruction to try and pay attention to their experiences the following occurs. They tacitly read this as an instruction to attend to the best candidates for being the qualia associated with their experiences, and decide whether *those very items* are properties of experiences themselves, or instead perceptible items like ordinary physical objects and/or their properties. If so perhaps we can gloss transparency as follows. When we try to attend to the 'qualia' associated with our experiences, the items we attend to which would be[25] the qualia associated with the experiences we are introspecting if any items were, are

[24] I.e.successful perceptual, illusory or hallucinatory.
[25] This is obviously to be read as indicative rather than counterfactual.

perceptible items, ordinary objects and/or their perceptible properties like colours and shapes. Depending on whether we take 'qualia' to by definition be properties of experiences (or other mental items), we might more simply phrase this as follows: when we try to attend to the qualia associated with our experiences, while introspecting, we discover that those qualia are in fact perceptible items, rather than properties of experiences or other mental items.

If we formulate the transparency thesis in this way, the problems created by the fact that we can, in some sense attend to properties of our perceptual experiences, when we introspect them, given that we certainly think about those experiences when doing so, go away. Those properties might be attended to, but if they are not the items which we recognize as being the items which would be taken by believers in qualia to be the qualia associated with the experience, then the fact that we attend to them when we introspect is just irrelevant to the truth of the transparency thesis, on the currently suggested formulation.

Formulating transparency in this way does imply that, at the very least, philosophers who have worked through the literature described in the previous Section, have enough sense of what is meant by 'qualia' to be able to know, *de re*, of certain items they attend to when they introspect, that those very items would be the qualia associated with their experiences, if any such items were. But this does not seem a huge leap beyond the basic claim that the text we discussed in the previous Section succeeded in fixing any sort of meaning for 'qualia' at all. Further, there are several things which make holding, that transparency is a claim about the sort of items we attend to, when we introspect our experiences, and attend to the items which are the best candidates to be their 'qualia', plausible.

Firstly, re-read the passages above from Harman, Tye and Martin, in which they defend the transparency thesis, but this time read instructions to try and attend to your experiences, or descriptions about what we attend to, when we try to attend to our experiences, as instructions

about attending to the 'qualia' associated with your experiences. I think the passages in question don't lose any persuasive force they otherwise have, when read in this way.

Secondly, many proponents of the transparency thesis have been happy to phrase their view as the view that 'qualia ain't in the head'[26], or in other words, have taken transparency to be the view that when we attend to the qualia associated with our experiences, we find them to be perceptible, rather than mental items. This suggests that they themselves think of themselves as making the point that the items which the philosophers in the previous Section thought were qualia, are in fact perceptible items and not properties of experiences (or other mental items).

Thirdly, and finally, we saw above that proponents of transparency have typically taken a relational view of phenomenal consciousness, or at least of phenomenal consciousness in the successful perceptual case. In particular they have been committed to the view that the phenomenal properties of our successful perceptions are fixed by the perceptible items, which, according to the transparency thesis (however exactly we formulate it) we would attend to were we to introspect those experiences. If the transparency thesis is really the claim that the best candidates to be qualia associated with our successful perceptual experiences are perceptible items, rather than experiences or their properties, then it's easy to see why those particular perceptible items, the ones which are the best candidates to be the qualia associated with our experiences, would be involved in constitutively fixing the phenomenal properties of our experiences. As we saw in the previous Section, there was always supposed to be some sort of very close and non-causal connection between the qualia associated with experiences, and their phenomenal properties. So an advantage of interpreting transparency in the way I'm proposing is that it makes sense of the use of the transparency thesis to support the claim that the phenomenal properties of successful perceptual experiences constitutively depend on certain perceptible items associated with those successful perceptions. Comparatively, it's not

[26] The slogan comes from Byrne and Tye (2006). See also Lycan (1995) for another account of qualia as perceptible items rather than properties of experiences.

immediately clear why we would think that the items we attend to or are directly aware of when we introspect our experiences should be involved in constitutively fixing their phenomenal properties, unless we are already tacitly assuming that the items we are directly aware of are qualia, or at least are the items which believers in qualia have taken to be qualia[27] [28].

So there is much to be said for defining transparency as the claim that, when we introspect our perceptual experiences and try to attend to the items associated with those experiences which proponents of the traditional internalist view of qualia would take to be their qualia, then those items we attend to which would be so taken, are perceptible items, and not experiences or their properties.

However it's cumbersome to continuously have to find some way of picking out, collectively, the items associated with all our successful perceptual experiences, such that believers in qualia in the traditional sense would, were they to undergo those experiences, take those items to be their associated qualia. An easier way, I think, to pick out those items, is to talk of the items an experience makes '**phenomenally present**' to its subject. It's been noted very often by philosophers of otherwise quite different background commitments, that there is some sense in which (visual) perceptual experiences seem to involve some sort of 'presence' 'within' the experience itself of some of the items which they count as perceptions of, and that nothing like this phenomenon need occur in the case of conscious thoughts. For accounts of this see Searle (1983), Sturgeon (2000), Martin (2002), Schroer (2012) and Frey (2013) amongst many others[29].

[27] Some might think it's analytic that qualia are properties of experiences.

[28] Though one might make an argument here from the fact that the phenomenal properties of our successful perceptual experiences seem to co-vary with the properties of the items we perceive by undergoing them.

[29] ...sensing is just different from...thinking. More exactly...sensory states involve a certain...apparent presence of an object which is simply not required in the case of pure thought...(Martin, 2002, pp.387-8).
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, p.9).
If, for example, I see a yellow station wagon in front of me, the experience I have is directly of the object. ...The experience has a kind of directness, immediacy, and involuntariness which is not shared by a belief I might have about the object in its absence…The visual experience I will say does not just represent the state of affairs perceived; rather,

Furthermore, I think that it is items which are 'present' to us in this sense which play the qualia role and fix the phenomenal properties of our experiences. I will therefore from now on refer to the items associated with perceptual experiences which constitutively fix their phenomenal properties, whose status as either perceptible items or mental items is relevant to the truth or falsity of the transparency thesis, as those which our experiences make phenomenally present to us. Those who don't think that the metaphor of 'presence' makes it any easier to glom on to the set of items being picked out can think of it as literal shorthand for 'items associated with an experience which are those which would be identified as best candidates to be 'qualia' in the traditional internalist sense by competent users of that term'. On my view though there's some independent value to bringing in the presence metaphor in clarifying just what items associated with experiences we are picking out here.

We can now formulate the transparency thesis as follows:

Transparency: our perceptual experiences, whether hallucinatory, successful perceptual or illusory never make phenomenally present to us themselves or any mental properties or other mental items; further, in the case of successful, non-hallucinatory or illusory, perceptions, they do make present to us either the objects they are perceptions of or some of the properties or tropes of those objects[30], or both.

---

when satisfied, it gives direct access to it, and in that sense it is a presentation of that state of affairs. (Searle, 1983, p. 46)

Although both the visual experience and the conscious thought make you aware of the station wagon in an unmediated manner, in the former the station wagon seems present before the mind in a way that it doesn't in the latter.' (Schroer, 2012, p.596), discussing the previous quotation from Searle.

In all sensory modalities, the objects of perceptual experience are there, present to us, in a way that the objects of most beliefs and judgments are not. (Frey, 2013, p.71).

Of these authors, Martin is a naive realist. Searle and Schroer, meanwhile are both representationalists, but Schroer is sympathetic to reductive causal accounts of representation itself whilst Searle is opposed to such accounts. Sturgeon's views are more difficult to classify. Frey does not put forward a positive view in (2013), but does deny representationalism.

[30] I'll generally leave tacit the possibility that it's perceptible tropes rather than perceptible properties our experiences make phenomenally present to us, except where it's relevant to the argument at hand.

Note that if someone disagrees with me, and holds that the transparency thesis is better formulated in terms of a notion of 'direct' attention, rather than a notion of attention to what is phenomenally present to us, this will not in fact matter much for the purposes of either the rest of this chapter, or the rest of this Thesis. The reason for this is that, to get a tension between the HMT and transparency, we only need it to be the case that there is a plausible formulation of the transparency thesis on which it entails that there are sometimes items associated with successful perceptual experiences which both fix the phenomenal properties of experiences, and which include ordinary physical objects and/or their perceptible properties.

For, as we'll see below, the tension between transparency and the HMT, comes about in the following way. Transparency entails[31] that the items associated with successful perceptions that constitutively fix their phenomenal properties are the sorts of things we perceive through undergoing successful perceptions, ordinary objects and/or their properties. But pure hallucinations by definition do not count as perceptions of any object. So we might worry that there is no non-arbitrary way of associating them with particular objects, and hence that there cannot be, for any pure hallucination, *H*, a particular object *o* associated with it, such that *o* plays a constitutive role in fixing *H*'s phenomenal properties. If particular objects are amongst the items associated with successful perceptions which play a constitutive role in fixing their phenomenal properties, there's therefore a worry that pure hallucinations will not be able to share all the phenomenal properties of successful perceptions. It doesn't actually matter for establishing this tension how exactly we pick out the items associated with successful perceptions which play a constitutive role in fixing their phenomenal properties. In particular it doesn't matter whether we pick them out as those experiences make available for direct attention, or those they make phenomenally present to their subjects. What matters is that the standard defences of transparency give us reason to think that it's items of the sort we perceive in the world which are

[31] Strictly speaking it entails this only given the further assumption-which I am making-that the items associated with an experience which fix its phenomenal properties are always items which that experience makes phenomenally present to its subject.

associated with successful perceptions and constitutively fix their phenomenal properties. More generally through this Thesis I give arguments which assume that the items associated with perceptual experiences which fix their phenomenal properties are those they make phenomenally present to their subjects. However, though I admit I cannot prove this rigorously, I suspect that, insofar as it's possible to define a notion of 'direct attention' which can be used to give an adequate statement of transparency, these arguments can be re-run with references to phenomenal presence being replaced with references to the items our experiences make available for the direct attention of their subjects without loss of plausibility.

We are now almost ready to move on to giving a precise account of the tension between transparency and the HMT. Before that, however, I want to say a little to introduce the two main types of theories of the metaphysics of phenomenal consciousness favoured by proponents of transparency. Doing so will then allow me to discuss in Sections Four and Five where various particular theories of these two general types stand on the question of whether we must give up the HMT to accommodate transparency.

One family of widely endorsed theories motivated by the transparency thesis are representationalist theories of consciousness. On these theories, part of what it is to be a phenomenally conscious experience is to represent the world as being a certain way[32], and the phenomenal properties of experiences depend (with metaphysical necessity) partly on their representational content. I say partly depend, because almost everyone agrees that there can be unconscious representation, and further many agree that in many cases exactly the same content can be represented consciously or unconsciously. If so there must be something other than their associated representational content which determines which content-bearers are phenomenally

[32] There are also possible versions of representationalism which apply only to a subclass of the category of phenomenal states. Dretske (1995, p.103) explicitly expresses doubts about whether his theory can account for certain emotional states which he thinks don't appear to bear representational content at all. (He gives 'joy, despair, anxiety and dread' as examples.)

conscious at all[33]. Representationalists hold, however, that for phenomenally conscious content-bearers, their phenomenal properties supervene on their representational properties, with metaphysical necessity[34]. Representationalism is often combined with the 'displaced perception theory of introspection', developed by Dretske (1995). On the displaced perception theory, when a subject *S* introspects a particular phenomenally conscious perceptual experience *E* and thereby comes to know what phenomenal properties *E* has things work like this. S first uses *E* to attend to the items in the world which it represents, and then automatically infers that she's undergoing the sort[35] of phenomenally conscious experience which represents the world as containing items just like those.

The other most popular type of transparency inspired theories of phenomenal consciousness are naïve realist theories. Prominent naïve realists include Fish (2009) and (especially) Martin (2002a, 2004, 2006, forthcoming). Campbell (2002) and Brewer (2008, 2011) are also often referred to as naive realists, but it is less clear that their views are, officially speaking, views about the metaphysics of phenomenal consciousness, in the sense of 'phenomenal consciousness' introduced above. Brewer does not use 'phenomenal consciousness' and its cognates, whilst Campbell's use of phenomenal vocabulary is somewhat idiosyncratic and differs from mine[36]. In the current context, we are only interested in naive realist theories insofar as they are theories about phenomenal consciousness, and so bear on the clash between transparency and the HMT. On naïve realist theories, at least when construed as theories of phenomenal consciousness, for a successful perceptual experience to have a particular phenomenal property *P*, is for it to bear some special relation, *R* to an object which bears some appropriate property *P*'. For example,

[33] See Chalmers (2010, pp.342-9) for more on this.
[34] At least, this is the strong form of the doctrine. There is also a weaker view on which the phenomenal properties of conscious experiences supervene on their representational content, plus their sense modality. See once again Chalmers (2010, p.347-8).
[35] If the phenomenal properties of conscious experiences supervene on their representational properties, then for each conscious experience, its representational content will determine some unique maximal phenomenal property.
[36] Martin does not always make use of phenomenal vocabulary either, but (forthcoming) makes clear that he conceives naive realism as a view about phenomenal consciousness.

according to naive realism, for a fully successful perceptual experience, *E*, to have the phenomenal property distinctive of experiences of say, red_{25}, it must bear *R* to some object *o* which *is* red_{25}. *R* is held by naive realists to only hold between experiences (or perceivers) and items in the world they count as successful perceptions[37]. Naïve realists disagree about what exactly to say about what illusory or hallucinatory experiences having phenomenal properties consists in, and even whether hallucinatory experiences have any phenomenal properties at all[38]. Transparency is then explained, at least for successful perceptual experiences, by the claim that when we introspect our experiences, we attend only to (all or some) of the items on the other end of the special naïve realist relation.

Two further points about naive realism and representationalism:

Firstly, although representationalist and naïve realist theories are by far the most popular sort of theories of the metaphysics of phenomenal consciousness which are designed primarily to accommodate transparency, they clearly do not exhaust the possible space of theories which do this. All a theory has to do to be consistent with transparency (as I have formulated it) is to hold that the items which are made phenomenally present to us by our successful perceptual experiences include only some or all of the sorts of items we perceive in the world, i.e. ordinary physical objects and/or the perceptible properties which ordinary physical objects can instantiate, such as shape and colour properties. In principle a transparency advocate could hold that there is a single, non-representational relation, *R*, such that: a) all phenomenally conscious perceptual experiences, whether successful perceptual, illusory or hallucinatory bear *R* to one or more *properties*, and b) which properties an experience bears *R* to plays a constitutive role in fixing its phenomenal properties.

[37] Naive realists agree that hallucinating an object and its properties never involves bearing *R* to anything, but they disagree about whether having an illusory experience of an object is a way of bearing *R* to that object.

[38] Fish (2010) explicitly commits to the startling view that they do not.

Secondly, although they are often presented as rival theories, and some naïve realists (e.g. Travis (2004), Brewer 2011)) deny the claim that phenomenally conscious perceptual experiences bear representational content at all, naïve realism and representationalism as typically defined are not, strictly speaking, inconsistent. Naive realism recall, is roughly the view that the having of phenomenal properties by veridical experiences is wholly or partially constituted by those experiences bearing some relation *R* to the items they count as successful perceptions of, a relation which no pure hallucination bears to any items. This is not strictly inconsistent with the view that the phenomenal properties of mental states supervene on their representational content, even though it entails that hallucinations, since they do not bear *R* to anything, cannot have exactly the same phenomenal properties as veridical experiences. Representationalists are not committed to the view that, for each veridical experience, there is some purely hallucinatory experience which bears the same content, since the content of experience might be singular in the successful perceptual case, and pure hallucinations might necessarily lack singular content. Indeed, there is nothing in the definition of naive realism which I gave above which entails that *R*, the naive realist relation borne by experiences to the objects they count as perceptions of, couldn't just be the *having as singular content* relation[39].

Section Four-Tension Between the HMT and Transparency

We can now outline where the potential for conflict between the HMT and transparency lies. On the one hand, we saw in Section One that it seems to be part of the meaning-fixing theoretical role of the term 'phenomenal properties', that which particular phenomenal properties an experience has is a matter of which qualia they have associated with them. On the other hand,

[39] Although obviously things are slightly more complicated on versions of naive realism on which *R* holds between subjects of experience, and perceived objects, rather than directly between experiences and objects.

we've seen in Section Two that if transparency holds, then the best candidates to be the items associated with successful perceptual experiences which play the role of qualia, and fix their phenomenal properties, are going to be the perceptible items, ordinary physical objects and/or their perceptible properties. For given transparency, these are the only items successful perceptions make phenomenally present to their subjects. But now worries of the following sort arise about the HMT. It is, to say the least, not obvious that pure hallucinations can make present to their subjects all the same kind of items as we can perceive in the successful perceptual case. In particular, there is reason to doubt that pure hallucinations can make physical objects themselves, as well as their properties, phenomenally present to their subjects. After all, it's not obvious that there is some non-arbitrary way to associate particular objects with particular hallucinations. Further support for this conclusion is arguably given (depending upon one's level of trust in the wisdom of philosophers!) by the fact that at most very few-quite possibly *no*-philosophers have ever defended the view that pure hallucinations can make particular objects phenomenally present to their subjects[40]. So we know from the first point about the meaning-fixing theoretical role of 'phenomenal properties', that experiences have their phenomenal properties in virtue of having particular qualia-role playing items associated with them. And we know from transparency that the qualia-role playing items associated with successful perceptions are the sorts of things we perceive in the world, physical objects and/or their perceptible properties. But now we've also seen reason to think that it's impossible for pure hallucinations to make particular physical objects phenomenally present to their subjects. Call any phenomenal property that can only be possessed by experiences which make objects phenomenally present to their subjects an '**object-requiring**' phenomenal property[41]. The worry is that, since the phenomenal properties of perceptual experiences depend on the items they make phenomenally

[40] I briefly defend this claim about the literature in sc.1 of ch.3.

[41] These properties are not necessarily object dependent in the sense that experiences can only have them because they make some particular object, *o*, phenomenally present to their subjects. Properties such that an experience cannot instantiate them unless it makes some object or other phenomenally present to its subject, but there is no specific object it must make so present, also count as 'object-requiring' by my lights.

present to their subjects, if successful perceptions make particular physical objects phenomenally present to their subjects, and pure hallucinations cannot do this, successful perception will instantiate object-requiring phenomenal properties, and pure hallucinations will not be able to instantiate such properties. This would obviously entail the falsity of the HMT.

Of course, a lot more would have to be said here in order to make a convincing case for the incompatibility of the HMT and transparency. We would need to know just how best to formulate the claim that experience's have their phenomenal properties in virtue of making particular items phenomenally present to their subjects. And we'd need to know that when formulated in this way, this claim, along with the truth about just what perceptible items successful perceptions make phenomenally present to their subjects, together entail that successful perceptual experiences instantiate object-requiring phenomenal properties. There are two obvious sorts of way in which this sort of argument for the claim that successful perceptions have object-requiring phenomenal properties might fail. Firstly, the perceptible items phenomenally present to us in the successful perceptual case might include properties, but not objects (or tropes[42]). Secondly suppose that objects are phenomenally present to us in the successful perceptual case. Still, if properties are also so present, perhaps the best account of exactly how the phenomenal properties of experiences constitutively depend on which items they make phenomenally present to their subjects will not entail that there are any object-requiring properties. Maybe, for example, it is only the properties which experiences make phenomenally present to their subjects that play a role in fixing their phenomenal properties, even though they make objects so present as well as properties. Or maybe, objects do play some kind of role in constitutively fixing the phenomenal properties of successful perceptions, but the properties in question are multiply realizable in some sense, and another way to possess each of them is to make some property or combination of properties or items built up from properties phenomenally present to a subject.

[42] The same considerations of arbitrariness which bear against the idea that it's possible for pure hallucinations to make particular objects phenomenally present to their subjects also, I think, bear against the idea that they can make tropes so present.

Still, we can see here the outline of a general strategy for arguing against the HMT, from transparency or stronger extensions of transparency. First, use transparency, or a stronger extension of it (or some other similar principle) to establish some condition, *C*, that all the qualia-role playing items associated with successful perceptions meet. In the case of transparency itself, this condition is just being perceptible items. Then show that pure hallucinations cannot make items which meet this condition phenomenally present to their subjects. Following on from this, use the fact that phenomenal properties are introduced as properties which experiences have in virtue of there being particular qualia-role playing items associated with them, to argue for some principle about what similarities there must be between two perceptual experiences, *E* and *E*', with respect to the items they make phenomenally present to their subjects, if *E* and *E*' are to match. And show that it follows from this constraint on matching, plus the fact that successful perceptions make items which meet *C* phenomenally present to their subjects, and pure hallucinations cannot, that pure hallucinations cannot match successful perceptions.

In the rest of this Thesis, I shall be interested in whether we can get an argument of this type against the HMT, involving, respectively, the transparency thesis, and a claim about what similarities in the items two experiences make phenomenally present to their subjects are needed for two experiences to share all their phenomenal properties which I call the '**Qualitative Requirement on Matching**' (hereafter **QRM)**. I'll properly introduced and defend the QRM soon, but for now we can think of it as being roughly the claim that if one experience is to match another phenomenally, there must be no qualitative difference in the items it makes phenomenally present to its subjects.. I explore this issue in ch.2 and 3, eventually coming down against us being able to maintain both the QRM and the HMT at the same time, given an extended version of the transparency thesis, which I shall argue should be endorsed by those who accept the original version of the transparency thesis stated in the previous Section. In ch.4, I argue that if we should drop one out of the QRM and the HMT, it is the HMT that should go. I won't discuss in this Thesis the third option of dropping transparency instead. This is not because

I think that we have vastly more reason to endorse transparency than the other two claims[43]. It's simply that with limited space, I am going to restrict myself to working out the consequences of transparency (given the assumption that it holds.)

Here is my official statement of the Qualitative Requirement on Matching:

Qualitative Requirement on Matching: Necessarily**,** two experiences, E_1 and E_2 can match phenomenally only if there is no qualitative difference in the items they make phenomenally present to their subjects.

So, what exactly is a 'qualitative' difference in the items which two experiences make phenomenally present to their subjects? And why should we think that, if there is a qualitative difference between the items which two perceptual experiences, E_1 and E_2 make phenomenally present to their subjects, then they must also differ in their phenomenal properties?

Well, recall that it's part of the meaning-fixing theoretical role of 'phenomenal property' that the phenomenal properties of experiences vary whenever their 'qualia' vary. Or at least, they were supposed to vary whenever they made qualia of a different *type* phenomenally present to their subjects. Here, recall, as we saw in Section Two, there was supposed to be a type of quale for each maximally determinate perceptible property, or certainly for perceptible colour and shape properties at any rate. For each of these perceptible properties, there was supposed to be a different-qualia type, such that all and only experiences which had a quale, *Q* of that type amongst their associated qualia would instantiate a particular phenomenal property. So, for example, for some particular minimally distinguishable shade of red, *R*, all successful perceptions of objects as *R* were supposed to have associated with them qualia of a particular type *Q*, unique to experiences of *R*. And so were illusory misperceptions of objects as *R* and hallucination as of

[43] In my view, the most serious problems for the transparency thesis are likely to be those it faces in accommodating empirical data about how attention shifts can change aspects of the perceived colour of a single object, apparently without there being misperception of its colour either before or after the change. See Block (2010), although Block officially frames his argument there as a challenge to naive realism and representationalism, rather than the transparency thesis itself.

objects that are *R*, And there was supposed to be a unique phenomenal property *P*, which all experiences which had *Q* as one of their associated qualia instantiated. Indeed, on the view on which 'qualia' were supposed to just *be* phenomenal properties, *Q* and *P* would simply be the same property, and hence identical. Now, as we saw when introducing the transparency thesis in Section Three above, believers in transparency, who think that phenomenal property or qualia talk is in anyway legitimate at all, have generally tried to get the perceptible objects which they think are phenomenally present to us to take over the role that qualia, in the traditional sense, were supposed to play in theorizing about phenomenal consciousness. The claim that two experiences differ in their phenomenal properties whenever there is a 'qualitative' difference in the items they make phenomenally present to their subjects is supposed to be what we get when we try to give an account of how the perceptible items which transparency claims are all that our perceptual experiences ever make phenomenally present to us, take over the function of 'qualia' in the traditional sense in fixing the phenomenal properties of our experiences.

In particular, just as there was supposed to be a qualia type, *Q*, for each perceptible property, under the old, internalist way of looking at things, if we are getting perceptible items to play the role of qualia, we can simply treat each perceptible property, as though it is a type of quale. In other words, whereas, under the old, internalist view, for each perceptible property there was a particular qualia-type, *Q*, and there was supposed to be a unique phenomenal property had by all and only those perceptible experiences that had a quale of type *Q* suitably associated with them, now that we are treating the perceptible items experiences make phenomenally present to us as playing the qualia role, we can simply say that for each perceptible property, *P*, there is some phenomenal property *P'* such that all and only perceptual experiences which make *P* phenomenally present to their subjects, are *P*'. It is differences in which properties experiences make phenomenally present to their subject that I am treating as 'qualitative'. The term 'qualitative' is being used here not to impart spurious plausibility to my account of matching via a pun on 'qualia', but simply to mark the fact that the condition the QRM lays down on matching

concerns the property-ish aspect of what experiences make phenomenally present to their subjects, rather than which objects they make so present. Admittedly, if we reject the existence of properties because of nominalistic scruples, the definition of a 'qualitative' difference here cannot be literally correct, since it mentions properties. If we are trope nominalists, we can easily fix things by saying that two experiences differ qualitatively whenever one makes present to its subject a trope of a property *P*, and the other makes no trope of *P* phenomenally present to its subject.

Things get a little trickier, however if we reject both tropes and properties, and think that objects alone are made phenomenally present to us by our successful perceptual experiences. A first thought might be that on such a view, two experiences differ qualitatively whenever one of them makes an object *o* that is *P* phenomenally present to its subject, and the other makes no object which is *P* phenomenally present to its subject. But this cannot be right, since not all ways that an object is are relevant to the phenomenal properties of any particular experience of it. For example, if an object *o* was once owned by Virginia Woolf, this does not mean that an experience *E* which makes *o* present to it subject must differ phenomenally from an experience, *E*' which makes an object *o*' which shares *o*'s shape and colour, but was never owned by Woolf, present to its subject. So it seems that if we reject the existence of both tropes and properties, we need a notion of an object being experienced as being one way, rather than another, to account for the fact that some properties of the objects phenomenally present to us are relevant to the phenomenal properties of our experiences and others are not. With such a notion in hand, we can then treat the QRM, on such a view, as the claim that two experiences *E* and *E*' differ phenomenally whenever one is an experience of an object, *o*, as *P* and the other is not an experience of any object as *P*. I won't at this point try to resolve whether such a nominalist should take experiencing an object as being some way to be require that that object actually is that way.

Note that I'm currently leaving open some questions here about the relationship between the objects, as opposed to properties, an experience makes phenomenally present to its subject, and its phenomenal properties, since the QRM just says that two experiences match only if they make the same properties phenomenally present to their subjects. It is silent on whether or not there are other requirements on matching. So it's silent on both of the following questions. Firstly it's silent on whether there are any whether there are phenomenal properties which experiences can only instantiate by making one particular object, *o*, phenomenally present to their subjects. I will assume that there are no such properties, since it seems that, in principle, successful perceptions of distinct objects can share all their phenomenal properties. A second and more difficult question is the following. Assuming that our successful perceptual experiences, at least, make both objects and properties (or objects and tropes) phenomenally present to their subjects, can two experiences match phenomenally despite making different numbers of objects phenomenally present to their subjects, provided that there is no qualitative difference in the properties (or tropes) they make present to their subjects? I will leave the answer to this question undecided for now. Somewhat arbitrarily, I will decide not to count such differences between experiences as 'qualitative'. But this is compatible with them entailing a phenomenal difference between experience, since the QRM is an 'only if' claim, not a biconditional.

There is, however one type of difference between perceptual experiences in the items experiences make present to their subjects beyond differences in which properties and relations they make so present which I will count as 'qualitative' and hence which my endorsement of QRM commits me to saying leads to a difference in the phenomenal properties of two experiences. More on this in ch.2, when I discuss the fact that our perceptual experiences can make different structured items, constructed out of the same properties and objects phenomenally present to us.

Section Five-Possible Responses to the Tension Between the HMT and Transparency

In the previous Section, we saw that there is a potential tension between the HMT and transparency for roughly the following reason. We know that the phenomenal properties of perceptual experiences are constitutively fixed by the qualia-role playing items associated with those perceptual experiences. And transparency tells us, in effect, that, in the case of successful perceptual experiences, the qualia-role playing items associated with those experiences are just the perceptible items they make phenomenally present to their subjects. However, we know that amongst the items we perceive in the world are particular physical objects, and that there are reasons to think that it's impossible for pure hallucinations to make particular objects phenomenally present to their subjects. The worry is then that, given that successful perceptions make objects phenomenally present to their subjects, and that the phenomenal properties of successful perceptions constitutively depend on the items they make phenomenally present to their subjects, successful perceptions might have object-requiring phenomenal properties. It might then prove impossible for pure hallucinations to instantiate such properties. We also saw that this tentative argument against the HMT from transparency, is an instance of a more general strategy for arguing against the HMT. First find a type of item, *T*, which successful perceptions make phenomenally present to their subjects. Show that it's metaphysically impossible for pure hallucinations to make items of type *T* phenomenally present to their subjects. And then show that this entails that pure hallucinations cannot match successful perceptions. To reach the final conclusion, evoke some principle about when a difference between two experiences, *E* and *E*', with respect to the items they make phenomenally present to their respective subjects, necessitates a phenomenal difference between those two experiences.

It turns out that many views in the literature can be seen as being shaped by the response to the challenge for those who want to hold on to the HMT raised by the fact that when we introspect

our phenomenally conscious experiences we seem to attend to perceptible items like ordinary objects and their properties which our experiences make phenomenally present to us. In this, the final Section of this chapter, I want to show how various well-known views in the literature on phenomenal consciousness, put forward by philosophers who accept transparency, can be classified according to how they respond to the tension between transparency and the HMT.

Most of the representationalists and naive realist views in the literature can be seen as reacting to this threat in one of three ways.

Firstly, some accept that the phenomenal properties of experiences are fixed by the items they make phenomenally present to their subjects, but hold that hallucinations make the same or relevantly similar items phenomenally present to their subjects, and so no conflict with the HMT arises. Here, relevantly similar is set by whatever principle(s) the philosophers accept about what similarities there must be in the items two experiences make present to their subjects before those items match phenomenally. The QRM plays this role for me.

For example, some philosophers hold that our successful perceptions make present to us properties, and properties alone, not particular objects[44]. This is consistent with transparency as I have formulated it, and as it understood by these philosophers, since they take the properties in question to be perceptible properties of physical objects like colours and shapes. These philosophers then go on to hold that whilst a pure hallucination, *H*, and a successful perception, *E*, only match phenomenally if they make the same items phenomenally present to their subjects, this is no threat to the HMT, since pure hallucinations can very well make properties phenomenally present to their subjects. More generally though, there is clearly a range of other possible views of a similar kind. These views accept some principle, whether the QRM or another, which specifies the similarities there must be in the items two experiences make present to their subjects, if those two experiences are to match phenomenally. But they go on to hold that this

[44] This view doesn't entail representationalism, but typically those who endorse it are representationalists.

does not generate problems for the HMT, because it is perfectly possible for pure hallucinations to be similar to pure hallucinations in the relevant way. For example, in ch.3, we'll meet a view of the following sort. Successful perceptions make structured items which have particular objects as parts phenomenally present to their subjects, and for two experiences to match phenomenally they both must make qualitatively identical such items phenomenally present to their subjects. But pure hallucinations can also make structured items which have particular objects as parts phenomenally present to their subjects. I call views of the sort discussed in this paragraph **Hallucinatory Presence views**. A Hallucinatory Presence view is any view of the following sort. Firstly, transparency holds. Secondly, in order for two experiences to match phenomenally, there must be some sort of similarity in the items they make phenomenally present to their respective subjects. But thirdly, the HMT comes out true anyway because it's possible for pure hallucinations to make the right sort of items phenomenally present to their subjects for them to match successful perceptions.

Against Hallucinatory Presence views we can set two other sorts of reaction which those who accept transparency can have to the threat of a clash between transparency and the HMT.

The first is simply to reject the HMT, on the grounds of simultaneously preserving transparency and some principle, as such as the QRM, which specifies what similarities there must be between two experiences in order for them to make the same items phenomenally present to their subjects. In other words, such philosophers hold that because there are certain types of item which (some or all) successful perceptions make phenomenally present to their subjects, but no pure hallucinations can, it's not the case that for every successful perceptions there a possible matching pure hallucination. I call such views **Phenomenal Difference** views. An important thing to remember is that to motivate such views, it's not enough to show that there's some type of item, ordinary physical objects say, which successful perceptions can make phenomenally present to their subjects, and pure hallucinations can't, in order to get the conclusion that the HMT is

false. You also need some assumption about what similarities in the items they each make phenomenally present to their respective subjects there have to be between two experiences if they are to match.

The second, meanwhile is to cling on to both transparency and the HMT, but simply reject the idea that for two experiences to match phenomenally, there must be some sort of commonality in what they make phenomenally present to their subjects. I call such views **Anti-Presence views**. A radical sort of Anti-Presence view, might hold that, even in the successful perceptual case the items our experiences make phenomenally present to us don't play any role in constitutively fixing the phenomenal properties of those experiences. Less radically, someone might hold that phenomenal properties are multiply realizable[45]: the items which are phenomenally present to us in the successful perceptual case do play a constitutive role in fixing their phenomenal properties. But for each phenomenal property that a successful perception can possess in virtue of making a particular item phenomenally present to its subject, there is another way in which experiences can possess that property. One might, for example, hold that there is some phenomenal property, *P*, which perceptual experiences can bear in virtue of making some particular physical object or other phenomenally present to their subjects, but that another way to bear *P* is to be an experience in virtue of undergoing which it *seems* to a subject that there is a particular physical object phenomenally present to her.

The distinction between Hallucinatory Presence, Phenomenal Difference, and Anti-Presence views is important because, as we'll see in the following chapters, there are differences in which types of objection the three views are open to. Since Hallucinatory Presence views hold that pure hallucinations can match successful perceptions because they can make items of the right type phenomenally present to their experiences, they are vulnerable to the objection that pure hallucinations can't really make items of the relevant type so present. But they do allow that the

[45] Thanks to Dave Chalmers for mentioning the possibility of this sort of view in conversation.

phenomenal properties of all types of perceptual experience, including both successful perceptions and pure hallucinations, are constitutively fixed by the items they make phenomenally present to their subjects. So they do at least allow us to preserve the claim that the phenomenal properties of perceptual experiences of all kinds are fixed by the qualia-role playing items associated with them. Anti-Presence views meanwhile don't entail any commitments about what is or isn't available for introspective attention in the hallucinatory case. But they do face the charge that, in detaching sameness of phenomenal properties from sameness of qualia-role playing items, they're simply changing the meaning of phenomenal vocabulary. Phenomenal Difference views meanwhile, are vulnerable to the charge that in rejecting the HMT, they are simply too counter-intuitive to be believed. In effect, then, the three types of view are open to very different sorts of objection: the first to at least somewhat theoretical and metaphysical objections about what could possibly be phenomenally present in the hallucinatory case, the second to objections about changing the subject which seem to be, in some sense of that vexed term, 'conceptual', and the third to charges of extreme counter-intuitiveness.

I will now show that many existing transparency inspired views in the literature fall into one of these three categories. Interestingly, it turns out that my three-fold classification cuts across the traditional divide between representationalists and naive realists/disjunctivists. Some possible naive realist views, as we'll see, are Anti-Presence views, whilst others are Phenomenal Difference views. And some representationalist views are clearly Hallucinatory Presence views, whilst others are Anti-Presence views.

Hallucinatory Presence Views:

A standard move amongst representationalists, as is pointed out by Pautz (2007), is to hold that even purely hallucinatory experiences make *properties*, as supposed to objects, phenomenally

present to their subjects. Such representationalists usually assume both that two experiences have the same phenomenal properties if and *only if*[46] they have the same existential representational content, and that all conscious perceptual experiences of whatever variety make phenomenally present to their subjects the items, or at least the properties, which enter into their representational content. Those items are presumed to include properties, but usually not objects, or at least not in the hallucinatory case. Their making those items phenomenally present to their subjects is then supposed to be constitutive of experiences having the phenomenal properties they in fact do. Representationalists of this sort usually emphasize that even in the pure hallucinatory case, there is some kind of genuine attention to or awareness of the properties perceptual experiences make phenomenally present to their subjects. In effect, the view is that the perceptible properties we represent play the qualia-role in both the successful perceptual and the purely hallucinatory case. This has seemed to many philosophers helpful in reconciling transparency with the HMT, since on the one hand, the phenomenal presence of perceptible properties in the successful perceptual case is compatible with transparency, and on the other it seems much less problematic to many philosophers to hold that properties can be phenomenally present in the purely hallucinatory case than to hold that objects can be.

Many representationalists have had views of this sort. For example, Dretske and Tye, the two main proponents of classical reductive representationalism both endorse this sort of view[47]:

> I am not saying that all experiences make one conscious of some *object.* The experiences one has during... hallucinations do not make one aware of any object; yet, these experiences are... as conscious as veridical experiences. As Millar (1991, p.2) puts it, "Nothing seems less problematic than that when we perceive something we are in a conscious state which we could be in even if we were not perceiving that thing."
>
> ...When I speak of experiences making a creature aware of something, I mean to be speaking of the properties, not the objects...the experience makes the creature aware of. In dreams and hallucinations, one occupies representational states that are not "tethered"….to an external object... In such cases, the internal representation "says" or "means" that something is pink and

[46] This if and only if claim is stronger than representationalism itself, which just requires the supervenience of the phenomenal on the representational across all possible worlds.

[47] Tye's view has changed in several respects since (1995), but it is still a Hallucinatory Presence view; see his (2014).

rat-shaped, but nothing…is pink and rat-shaped. …The person in whom this representation occurs is, to be sure, aware of these properties. They are exactly the properties (pinkness, rat-shapes, etc.) as those who see pink rats are aware of. …What makes a creature conscious of these properties is the same thing that makes a person who sees pink rats conscious of them: an internal state that represents something to be pink and rat-shaped.' (Dretske, 1995, pp.101-2).

As you view the scene before your eyes and how things look to you, necessarily, if any of the qualities of which you are *directly* aware change, then the phenomenal character of your experience changes. Consider, for example, the facing surface of a ripe tomato... In attending to the colour of the tomato, you are directly aware of a certain quality, *Q*, as covering that surface. …But change *Q*-for example, by changing the colour of the tomato or by donning colour-inverting lenses-and what it is like for you in viewing the tomato necessarily changes. Facts like this one are surely not brute. Moreover, they obtain even in the case that you are hallucinating. If the tomato does not exist, still you are directly aware of *Q*; and if some other quality replaces *Q*, the phenomenal character of your experience changes. An explanation is needed of why the phenomenal character of visual experiences is sensitive in this way to surface qualities of the surface experienced. …the explanation surely is that the phenomenal character *involves* the surface qualities of which the subject of the visual experience is directly aware-that these qualities at least partly *constitute* phenomenal character. (Tye, 2000, p.48.)

Here Tye and Dretske are endorsing the following view: hallucinatory experiences and successful perceptual experiences can, in principle make the same properties phenomenally present to their subjects, since in both cases, the properties they make so available are simply those which enter into the experience's representational content. And this is supposed to be enough to show that it's possible for there to be pure hallucinations which share all and only the phenomenal properties of successful perceptions[48].

Not all Hallucinatory Presence views are representationalist views though, or involve the claim that the items our experiences make phenomenally present to us are those which enter into their representational content. Mark Johnston (2004) has defended a view on which successful perceptions and pure hallucinations can in some sense be experiences of the same type, because they can make exactly the same complex properties phenomenally present to their respective subjects. But Johnston doesn't think that an experiences making a complex property

[48] They don't use my technical term 'phenomenally present' themselves, but it's clear that they are talking about the items associated with experiences which are the best candidates for being the qualia so associated, and which play a constitutive role in fixing the phenomenal properties of those experiences. Dretske does not use the term 'phenomenal property', but it's clear from the general context of (1995) that when he says that experience types are the same across the two cases, he means to say something which at least entails sameness of phenomenal properties between the relevant pairs of successful perceptions and hallucinations which share the same existential content.

phenomenally present to its subject is a matter of it having representational content into which that property enters. Indeed, he doesn't, so far as I can see, try to give any substantive account of what makes it the case that a particular experience makes a particular complex property phenomenally present to its subject. Assuming that Johnston thinks that pure hallucinations which make the same complex properties phenomenally present to their subjects match them phenomenally, and match them because they make the same complex properties so present, this would count as a Hallucinatory Presence view. I think Johnston is indeed best read as holding that if two experiences make the same complex properties phenomenally present to their subjects, this is indeed enough to guarantee that those experiences will match. In any case, it's of interest that there is a possible non-representationalist Hallucinatory Presence view of this sort, whether or not Johnston endorses it.

Anti-Presence Views:

At the same time as many representationalists have stressed the idea that *properties* can be phenomenally present to us when we undergo pure hallucinations, many other representationalists have instead stressed the idea that we can aptly describe pure hallucinations as being 'of' certain *objects* even though they make no such objects phenomenally present to their experiences. After all, other representational states, such as thoughts can undoubtedly be appropriately described as being about certain objects, even though those objects do not in fact exist. If we are representationalists, we might therefore think that perceptual experiences can be appropriately described as being 'of' objects which fit their representational content, whether or not any such object in fact exists. So, for example, Harman (1990) notes that one can, for example correctly describe someone as thinking about the fountain of youth, without committing oneself to the existence of such an object. According to Harman this is the correct model for thinking about perceptual experiences. Whenever an experience represents that there exists some object *o* with a property *P*, we can appropriately describe it as an experience of an object *o* that's *P*, as

an experience of a red square say, whether or not there in fact exists an object *o* which is *P* and which is the object the experience is an experience of.

Now, although Harman himself does not commit to the legitimacy of talk of phenomenal properties, there's a possible Anti-Presence view suggested by his remarks. On this view the phenomenal properties of experiences supervene on their representational properties; to have, say, the phenomenal property, *P* associated with successful perceptions of red_{47}, is to be a conscious perceptual experience that represents that there is some object which is red_{47}, and hence to be aptly describable as an experience of an object that is red_{47}. For an experience to be 'of' an object that is *P* though, on this view, doesn't entail that there is some object that is *P* which such experiences are experiences of. Rather talk of what items an experience is 'of' creates a context from which existence claims are not exportable. To use a concrete example, a successful perception, *E* of a red apple, and a pure hallucination, *H* of a red apple may share all their phenomenal properties because they both represent that there are apples of exactly the same shade of red, *R* and vaguely spherical shape *S* at a particular distance and direction from their subjects[49]. But there is no apple that is red and vaguely spherical which is the apple which *H* is an experience of or which *H* makes phenomenally present to its subject. This view is an anti-presence view because it characterizes what it is for experiences to share their phenomenal properties in terms that don't involve the idea that there are relevant similar items that they make present to their subjects. Instead, it says that they have to be appropriately describable as being 'of' relevantly similar objects, but for an experience, *E* to be 'of' an object isn't really for there to be some object, *o*, its related to in an appropriate way, and hence in particular does not consist in it making some object phenomenally present to its subject. Since the view holds that sharing phenomenal properties isn't a matter of making relevantly similar items phenomenally present to a subject, the view is an Anti-Presence one. From now on, I will refer to Anti-Presence

[49] The example is just for illustrative purposes, I'm not actually claiming that natural properties like *being an apple* enter into the content of our perceptual experiences, although see Siegel (2010, ch.4) for an argument that they are.

views of this type as 'intentionalist' views, since that is what such views were originally named in the 1950s.[50] Note that it might be possible to have a parallel view on which the phenomenal properties of experiences depended on which properties they represented to be instantiated in the world but, on which, strictly speaking there necessarily are no properties and so no successful perceptions or pure hallucinations in fact make properties phenomenally present to their subjects. I know of no such view being defended in the literature however.

Apart from representationalist anti-presence views there are also possible anti-presence views which can be characterized as naive realist. In particular, some remarks of Martin's in (2004) suggest such a view, although Martin himself does not clearly endorse it. In (2004) Martin claims that there is no type of mental state in common between a successful perceptual experience, and a hallucination which involves the same underlying narrow brain state, except that the hallucination has the property of being indistinguishable by introspection alone from the corresponding successful perceptual experience which shares the same proximal cause within the brain. Given that phenomenal properties are mental properties, this clearly forces us either to give up the idea that such a pair of experiences could share all their phenomenal properties-that is, give up on the HMT-or to hold that phenomenal properties just are properties of being introspectively indistinguishable from particular types of successful perceptions. Such properties are usually referred to as 'negative epistemic properties' in the literature. If we do identify phenomenal properties with negative epistemic properties of the sort described, then this will perception *E* there is some pure hallucination *H* which is introspectively indistinguishable from successful perceptions of all the same types as the types of successful perceptions which *E* is indistinguishable from, the HMT will be preserved. Further, having a negative epistemic property is not a matter of making some item present to a subject, so if the HMT does come out true on it, the resulting view will count as an Anti-Presence view. For it is a view on which the HMT is true, and on which two experiences matching isn't a matter of some similarity in what they make

[50] See Martin (forthcoming, ch.1) for some discussion of the history of the view.

realism, on my definition of the latter. Naïve realism (as I've defined it) might look at first glance to be incompatible with the claim that phenomenal properties are negative epistemic properties, since naïve realism says that successful perceptions bear their phenomenal properties in virtue of bearing some special relation *R* to ordinary physical objects. The property *P* of bearing *R* to some object *o* is obviously not a negative epistemic property. But in fact, this is not a problem so long as having properties like *P* is a way for an experience to have a negative epistemic property. And given that the naïve realist relation *R* is, recall, a relation which experiences bear to objects only when they are successful perceptions of those objects, having properties like *P* will at least entail that an experience *E* has certain negative epistemic properties. For since *E* bears *R* to some object it will be a successful perception. *E* will therefore be indistinguishable from all the types of successful perception which it itself is, so it will have some negative epistemic properties of the relevant type. These can presumably be identified with the phenomenal properties which it has in virtue of bearing *R* to *o*.

Now, as a matter of fact, it is fairly clear from (2004) that Martin does not identify phenomenal properties with such negative epistemic properties. Here is the key passage:

> Experiences which share the proximate causal conditions [sic] as veridical perceptions but which lack the further conditions [sic] for being perceptions cannot have Naïve phenomenal properties. But, as we shall see, there is reason to think that whatever kind of event these experiences are, the...same kind of event occurs when one is perceiving. (Martin, 2004, p.53).

Here 'naive phenomenal properties' are, presumably 'phenomenal properties', if phenomenal properties are as naive realists take them to be. Martin self-identifies as a naive realist. So he seems here to be denying that phenomenal properties really are the sort of negative epistemic properties which can be possessed by both successful perceptions and pure hallucinations that were mentioned in the previous paragraph. Presumably Martin's view is instead that all successful perceptions bear phenomenal properties which it's impossible for pure hallucinations to instantiate, but nonetheless someone undergoing a particular pure hallucination is, in some other

sense, 'the same kind of event' as someone undergoing a subjectively indistinguishable successful perception.

Phenomenal Difference Views

It seems likely that Martin's official view in (2004), meanwhile, would count as some sort of (naive realist) Phenomenal Difference view. For if we don't identify the 'naive phenomenal properties' Martin speaks of with negative epistemic properties, it's hard to see what else we could identify them with which would leave the view as a naive realist one except properties that experiences have in virtue of bearing the naive realist relation to particular objects. This is particularly so since Martin is committed to an account of phenomenal properties which entails that successful perceptions bear phenomenal properties which pure hallucinations with the same proximal causes cannot. It's hard to think of a principled view that will have this result without entailing that pure hallucinations more generally cannot share those properties. And Martin certainly endorses transparency (2002a), so if his view also entails the falsity of the HMT, it will count as a Phenomenal Difference view in my sense.

Unsurprisingly, Martin is not the only naive realist who endorses a Phenomenal Difference view. If one thinks that there are phenomenal properties which successful perceptions bear in virtue of their bearing the naive realist relation to objects (or perhaps objects and their properties) then it's natural to at least consider the view that bearing the naive realist relation to some appropriate object is, for each of these phenomenal properties, the *only* way perceptual experiences more generally can come to possess them. From the latter stronger view, the falsity of the HMT does follow, given that the naive realist relation is supposed to be a relation which experiences can bear only to items they are successful perceptions of. The resulting naive realist

view will be a Phenomenal Difference view, assuming that it is compatible with transparency, as naive realist views are intended by their proponents to be, and is a view on which the objects on the other end of the naive realist relation in the successful perceptual case are thereby made phenomenally present to us.

The naive realist who most clearly endorses a Phenomenal Difference view is Fish (2009). He endorses the radical claim that pure hallucinations, unlike successful perceptions, lack phenomenal properties altogether. He endorses this radical claim because he thinks that all phenomenal properties are properties experiences have in virtue of bearing the naive realist relation to facts, where facts are structured items which have objects as constituents. Facts on his view, and some of their constituents, are what is phenomenally present to us in the successful perceptual case. Pure hallucinations, on his view, don't bear the relevant relation to facts. (Fish's view is also a view on which transparency holds.) The attitudes of the two other naive realists I mentioned, in Section Three, Brewer and Campbell, are a little more difficult to determine. I think that, ultimately Campbell would endorse a Phenomenal Difference view, whilst it is difficult to determine whether Brewer, if he accepted the legitimacy of 'phenomenal property' talk at all, would hold to a Anti-Presence View or a Phenomenal Difference view.

Briefly, on Brewer (I omit discussion of Campbell for reasons of space):

Brewer does not himself use the term 'phenomenal property', either in his (2008) paper outlining his views on illusion and hallucination, or his (2011) book defending his total view of the nature of perception. (2008) is mainly an account of Brewer's views on perceptual illusions, and only mention hallucinations in passing. Brewer takes a similar line to Martin on hallucinations, holding that 'hallucinatory experiences have to be characterized by giving a qualitative description of a mind-independent scene, and saying that the subject is having an experience which she cannot distinguish by introspection alone from one in which the constituents of such a scene are the direct objects' (2008, p.174). This is enough to tell us that Brewer probably doesn't accept a

Hallucinatory Presence view, since his characterization of hallucinations makes no reference to them making any items phenomenally present to their subjects, or more generally, to the idea that they involve any relation holding between experiences and other items at all. However, it doesn't tell us anything about whether Brewer thinks that the negative epistemic properties shared between successful perceptions and the hallucinations they are indistinguishable from, *are* phenomenal properties. Nor does it tell us whether he thinks that a successful perceptual experience *E*, and a subjectively indistinguishable hallucination *H* share any other mental properties in common with which phenomenal properties could be identified. So in (2008), Brewer does not commit directly to either a Phenomenal Difference or an Anti-Presence view. And nor does he say anything which obviously entails a view of either of these sorts. In (2011), things are much the same. Brewer defends the view that 'any given perceptual experience is most fundamentally to be characterized, *as the specific conscious condition that it is*, as a case of conscious acquaintance, from a point of view, in a particular sense modality, and in certain circumstances, with particular mind-independent physical objects'. And he also endorses the view that there is one salient way of typing experiences on which a pure hallucination, *H* is an experience of the same type, as a successful perception, *E* if *H* is subjectively indiscriminable from *E*. But he does not say whether such hallucinations can share (all and only) the *phenomenal* properties of the successful perceptions they are indiscriminable from[51]. However, his claim that the 'specific conscious condition' involved in a perceptual experience is 'fundamentally to be characterized' as involving a non-representational relation holding between the subject of the experience, his point of view and the perceived objects, suggests he might identify phenomenal properties with relational properties of perceptual experiences, or at least might think that successful perceptions bear their phenomenal properties in virtue of being appropriately related

[51] Brewer does use the term 'phenomenal character' once in (2011), but only in a discussion of whether successful perceptions of distinct but qualitatively identical objects can possess the same phenomenal character (p.101). His answer to this question is that 'phenomenal character' is ambiguous between readings on which these is and isn't possible, but he doesn't provide a full account of how successful perceptions gain their phenomenal character. Nor does he tell us directly whether he thinks pure hallucinations share the 'phenomenal character' of successful perceptions.

to the objects they count as perceptions of. For example, perhaps he would endorse the claim that a perceptual experience, *E*, bears some phenomenal property *P* iff some relation, *R*, holds between *E*, some appropriate objects, and a point of view, where *R* is a relation which can only hold between an experience and an object if the experience counts as a perception of that object. If so, his view will be a phenomenal difference view, since each successful perception will have phenomenal properties of this kind which any metaphysically possible pure hallucination will lack. But it is difficult to be sure when Brewer does not commit himself explicitly.

Chapter Two: Against Property-Based Hallucinatory Presence Views

All things being equal, we should want to simultaneously preserve transparency, the HMT, and the QRM. The QRM seems to be something like a conceptual truth (in some sense of that vexed term). I find the intuitive case made for transparency by authors like Tye and Martin very

compelling, as do many others. And denial of the HMT seems very counter-intuitive[52]. A view which did this, would of course, count as a Hallucinatory Presence view in the classificatory scheme of the last chapter. In this chapter, and the next, I argue that, sadly, we cannot preserve all of the HMT, the QRM, and the transparency thesis, because all Hallucinatory Presence views are false. My argument for this will run as follows.

Firstly, in this chapter, I draw attention to the fact that our successful perceptions make phenomenally present to us not mere pluralities of single items, but rather structured items. I show that two experiences can differ phenomenally because of differences in the structures they make phenomenally present to their subjects. This, I will argue, should cause anyone who endorses the QRM, to endorse a claim I make about which types of differences in the structured items two experiences make phenomenally present to their respective subjects entail a difference in the phenomenal properties of those experiences. I call the resulting principle the extended QRM. I then go on to argue that we should also endorse a particular extension of the transparency thesis. This extended version of transparency entails that structured items which have particular objects as parts are phenomenally present to us in the successful perceptual case; these items I call 'perceptible scenes'. From this conclusion, plus the extended QRM, it follows that possible pure hallucinations will not match successful perceptions unless they too make perceptible scenes phenomenally present to their subjects. From this claim about pure hallucinations it in turn follows that a popular subset of Hallucinatory Presence views are false: those on which pure hallucinations can match successful perceptions, despite not making objects or structured items that include objects as parts, phenomenally present to their subject, because they instead make properties and items built out of properties so present. In the rest of this chapter, I will refer to such views as 'property-based Hallucinatory Presence views'.

[52] In ch.4 I will argue that our reasons to endorse the HMT are not as strong as we might initially think, but I still think that point about counter-intuitiveness carries some wait.

Having established in this chapter that pure hallucinations must make structured items which have ordinary objects as parts, phenomenally present to their subjects if they are to match successful perceptions, I argue in the next chapter that no possible pure hallucination makes any such item phenomenally present to its subject. This shows that if I am right that those who endorse the QRM should also endorse the extended QRM, and that those who endorse transparency should endorse extended transparency, we cannot simultaneously preserve transparency, the QRM and the HMT. More generally, if we cannot jointly accept all three claims I think that this shows that all Hallucinatory Presence views are false. Hallucinatory Presence views, recall, are supposed to be views on which we can reconcile the following. Firstly transparency and the HMT. And then whatever the correct principle is about what similarities there must be between two experiences with respect to the items they make phenomenally present to their subjects in order for those two experiences to match phenomenally. I've already argued in the previous chapter that the QRM should be at least part of or entailed by this principle[53]. So if I am right that we cannot simultaneously hold onto the QRM, the HMT, and transparency, then this refutes Hallucinatory Presence views.

This chapter has three Sections.

In the first, I explore one very quick argument against Hallucinatory Presence views, and explain why I'm not happy with it. In the literature on naive realism, it's often claimed that it's just obvious that our successful perceptual experiences make particular physical objects, and not just perceptible properties, phenomenally present to us. From this, it might be thought, it follows that successful perceptions have object-requiring phenomenal properties, since the phenomenal

[53] One could admit that, all things being equal, it is best if a view preserves the QRM, but try and come up with a similar but weaker constraint on matching that is consistent with the HMT and transparency. Such a view might be then claimed to be superior to one on which we either had to give up one of the latter principles, or we had to drop the idea that the phenomenal properties of experiences are constitutively fixed by the items they make phenomenally present altogether. But I think that there is no natural weakening of the QRM which both looks reasonably principled and non-gerrymandered, and can also account for the cases in which everyone will agree that two successful perceptions differ in their phenomenal properties despite being perceptions of the same atomic properties and objects.

properties of successful perceptions are constitutively fixed by the items they make phenomenally present to their subjects. If successful perceptions really do have object-requiring phenomenal properties, it would follow from this, plus the widely accepted claim that it's impossible for pure hallucinations to make particular objects phenomenally present to their subjects that the HMT is false, and hence that so are Hallucinatory Presence views. Something like this argument seems to lurk in the literature on naive realism, though it's not clear if anyone actually endorses it. In Section One, I explain why I don't find it convincing.

In the second, I specify and explain the sense in which I take our experiences to be of structured items and argue for the extended version of the QRM.

In Section Three, I give my intuitive case of thinking that successful perceptions make structured items, so-called 'perceptible scenes', which have ordinary objects as parts phenomenally present to their subjects. I also show that it follows from this extension of the transparency thesis, plus the extended QRM, that if pure hallucinations cannot make perceptible scenes phenomenally present to their subjects, no possible pure hallucination matches a successful perception.

Note that I am not claiming the argument I give here *proves* that property-based Hallucinatory Presence views are false. There are varying points at which the argument can be resisted which I will flag below. My claim is merely that my argument provides some, reasonably strong, but not necessarily completely compelling reasons to doubt such views.

Section One-Previous Attacks on Property-Based Hallucinatory Availability Views

In the literature it is quite common to find discussions of the claim that when we introspect our successful perceptual experiences, introspection reveals that they make particular physical

objects and not just properties phenomenally present to us. This claim appears especially frequently in discussion of naive realism, where naive realism is, roughly, the view that there are interesting and substantive properties of successful perceptions which they have in virtue of bearing some special relation to the objects they count as perceptions of. It is assumed that pure hallucinations cannot instantiate such properties, roughly because they cannot make objects phenomenally present to their subjects[54]. Now suppose we take the important class of properties in question to be phenomenal properties[55]. Then this suggests the following argument against property-based Hallucinatory Presence views. Successful perceptions make ordinary physical objects phenomenally present to their subjects. If a perceptual experience, *E* makes some ordinary physical object phenomenally present to its subject, then no perceptual experience *E'* which does not make particular objects phenomenally present to its subject can match *E*. Therefore pure hallucinations cannot match successful perceptions just by making properties, or items built up from properties phenomenally present to their subjects. Call this the Quick Introspective Argument or QIA for short.

Now, it's not clear in fact, just how widely endorsed the QIA is. Although naive realists often seem to endorse the claim that successful perceptions have object-requiring phenomenal properties and often claim that it's introspectively obvious that successful perceptions make ordinary objects phenomenally present to their subjects, they do not generally argue directly from the latter claim to the former. Instead, at least as Martin[56] and Fish (2009)[57] present things,

[54] See for example the survey article Nudds (2009): 'it is widely accepted that experiences seem to have the NR [i.e. naive realist] property' (pp.334-5)', where to have the naive realist property is supposed to be having particular objects as 'constituents'.

[55] Fish (2009) begins his defence of naive realism with a discussion of the claim that our experience of ordinary objects (and their properties) is 'direct'. Later (pp.7-9) Fish makes clear that he thinks the objects we directly experience play a constitutive role in fixing the phenomenal properties of the successful perceptions through which we experience them.

[56] See Martin (2002a) for the claim that Anti-Presence views of the sort I called 'intentionalist' in the last Section, can capture the introspective data involving the apparent presence of objects in the successful perceptual case just as well as his own naive realist view. It is possible however, that Martin merely concedes this for the sake of the argument in (2002a) as in (2006) he gives the following as the motivation behind his 'disjunctivist' version of naive realism: 'the disjunctivist advocates Naive Realism because he or

there are two views which can equally well accommodate the fact that objects seem to be phenomenally present to us in the successful perceptual case: firstly some sort of naive realist view on which experiences have object dependent phenomenal properties. And secondly the sort of 'intentionalist' Anti-Presence view I discussed in ch.1. They do not directly discuss, on the other hand, property-based Hallucinatory Presence views, seemingly assuming that the opponents of naive realists will either endorse intentionalism, or adopt some sense-datum or old-style internalist qualia view incompatible with transparency. This makes it hard to know whether they would endorse the claim that property-based Hallucinatory Presence views are incompatible with successful perceptions making objects phenomenally present to their subjects.

Regardless, what should matter to us here, is whether it provides us with a good argument against property-based Hallucinatory Presence views. I think that it would be foolish to rely on the QIA, for the following two reasons:

1) Firstly there is an underappreciated problem about how we get from the conclusion that objects are phenomenally present in the successful perceptual case, to the conclusion that successful perceptions have object-dependent phenomenal properties. We saw in ch.1 that there is a strong case for thinking that differences in the *properties* two experiences make phenomenally present to their subjects, entail phenomenal differences between them. But we also saw there that it's much less clear that there any *object*-requiring phenomenal properties. This is because, recall, properties look to be the best candidates for playing the role of qualia, given that there are no qualia, in the traditional, internalist sense of the term, and it is part of the meaning fixing theoretical role of 'phenomenal property' that the phenomenal properties of two experiences vary when

she thinks that this position best articulates how sensory experience seems to us to be just through reflection' (p.354).

[57] See (2009, ch.1, pp.21-22) for Fish's claim that the sort of 'intentionalist' Anti-Presence view discussed in ch.1 can explain the apparent phenomenal presence of objects in the successful perceptual case as well as the naive realist can.

the ‘qualia’ associated with them vary. In comparison, we saw that nothing that was originally said about what phenomenal properties were supposed to be when the term was introduced seemed to guarantee that if one experience makes *objects* (or a particular object) present to its subject and another does not, those experiences should differ phenomenally.

2) There is also a problem of disagreement for the naive realist. Presumably those philosophers who endorse property-based Hallucinatory Presence views, are not convinced by the QIA. And many philosophers do endorse such views: in the previous chapter, we saw that for starters, Dretske (1995), Tye and Johnston all do and Pautz (2007) gives an extensive list of others. All of the contemporary proponents of property-based Hallucinatory Presence views are presumably aware of the idea that introspection reveals the phenomenal presence of objects in the successful perceptual case, and that this tells against their view that pure hallucinations can match successful perceptions by making properties (or items built up from properties alone) phenomenally present to their subjects. Now, there are two grounds on which one can reject the QIA. One might reject it for the reasons given under 1), namely that the move from the claim that objects are phenomenally present in the successful perceptual case to the claim that there are object dependent phenomenal properties is dubious. But one might also simply reject the notion that we can tell, just by introspection that our experiences make objects, and not just properties phenomenally present to us. Insofar as Tye, for example, extensively discusses what exactly we find when we introspect, in order to defend transparency, the following is notable. He doesn’t concede that objects are phenomenally present to us in the successful perceptual case, and then go on to explicitly reject the claim that this proves that successful perceptions instantiate object-requiring phenomenal properties. Nor does he even concede that it *seems* as though successful perceptions make objects as well as properties phenomenally present to us at first glance, and try to explain this away.

Instead, he simply assumes that holding that perceptible properties of the sort which ordinary objects instantiate are present in the successful perceptual case is all a view needs to do to respect the introspective evidence which motivates transparency. This strongly suggests that Tye at least does not find it introspectively obvious that some of the items phenomenally present to us in the successful perceptual case are objects and not properties. Furthermore, it's hard to find any proponent of a property-based Hallucinatory Presence view who concedes this, or even concedes that it seems to be true at first glance[58].

Now, insofar as there are philosophers who have mastered the relevant literature on perception and who have presumably considered the matter carefully who don't think that we can distinguish introspectively between an experience making an object phenomenally present to us, and its making a property so present, this causes difficulties for the QIA.

In particular it casts doubt on the claim that objects are present to us in the successful perceptual case, at least if the only evidence we're supposed to have for that claim is that when we introspect it seems as though objects are so present. Now, we have to be careful about several senses of 'cast doubt' on that are possibly in play here. Most ambitiously, we might think that the disagreement here ought to cause even those philosophers who do find it introspectively obvious that objects, not just properties are phenomenally present in the successful perceptual case to abandon the claim that they are so present, or at least drastically reduce their credence in it. I would not want to endorse anything this strong however. Issues about when disagreement between philosophers about what is intuitive (or in this case, introspectively obvious) ought to cause partisans in the dispute to abandon, or even reduce their confidence in their claims

[58] Pautz (2007) explicitly denies that we can distinguish introspectively between objects being phenomenally present to us and properties being so present. He himself doesn't endorse any kind of Hallucinatory Presence view, but his denial is no less evidence against the claim that it's just obvious that objects are phenomenally present to us in the successful perceptual case for all that.

about the matter of disagreement are *extremely* complex[59]. A weaker, and more plausible claim though is that, if there is widespread disagreement here about what introspection actually reveals, this at the very least greatly reduces the dialectical effectiveness of arguments that depend on controversial claims about what introspection reveals that are not necessarily accepted by those holding the views the arguments target. An argument against property-based Hallucinatory Presence views that did not depend on any such controversial claim about what introspection reveals hence looks dialectically more effective, all things being equal, than one which does. So this gives us some grounds for preferring my argument, given in the next two Sections against such views, if, unlike the QIA it doesn't depend on any such claim.

Something like the second of these problems also seems to infect an argument given against property-based Hallucinatory Presence views by Pautz (2007). Pautz claims that it's *a priori* that nothing a visual experience makes phenomenally present to its subject can fail to be spatially extended[60] (2007, p.517). And he goes on to suggest that, for every sensory modality, there will be some other type of perceptible property, *T* such that it's a *priori* that perceptual experiences of that sensory modality will make objects with properties of type, *T*, not just properties of that type itself, phenomenally present to their subjects (2007, p.517)).

But if these claims are really so obviously *a priori*, it's not clear why they are not widely acknowledged. After all, they are quite simple and general, so one would think that if they were obvious *a priori* truths, then most philosophers thinking about perception would be aware of them. Yet it's hard to find examples of proponents of property-based Hallucinatory Presence views admitting that such principles seem compelling and trying to explain them away, or of other opponents of such views relying on them. Further, Pautz does not give any further

[59] See Williamson (2007, ch.7) for some of the complexities here.

[60] Actually Pautz phrases the point in terms of 'visual awareness' rather than phenomenal presence, but I think he would accept the translation into my terms, and in any case, what matters for my current purposes is the plausibility of the claim when phrased in terms of 'phenomenal presence'.

argument for the principles, so he can hardly claim that they are *a priori* but not obviously so. *Non*-obviously *a priori* claims require some argument to establish them.

Section Two-The Structured Items Phenomenally Present to Us

Nonetheless, I think there is still something plausible about the idea that views on which only properties and not particular objects are phenomenally present to us in the successful perceptual case, are somehow in tension with transparency. More precisely, I think there is something to be said for them being in tension with what I take to be the basic idea behind the transparency thesis . This is that there is no gap between successfully perceiving items in the world, and those items being present to one. Here, I try and work out a slower and more careful argument against property-based Hallucinatory Presence views, which starts from the idea that what is phenomenally present to us in the successful perceptual case is the totality of what we perceive in the world, not just the properties we perceive alone.

In particular, I've so far been working with the simplified assumption that what our experiences make phenomenally present to us are individual objects, or individual properties, or both. But in fact, what is phenomenally present to us in the successful perceptual case is much more richly structured than a mere plurality of individual objects or properties. Think, for example, of the difference between E_1, a successful perception of a red square, slightly to the left of a yellow rectangle, slightly to the left of a blue circle, slightly to the left of a green triangle, and E_2, a successful perception of a red square, slightly to the left of a blue circle, slightly to the left of a yellow rectangle, slight to the left of a green triangle. These experiences clearly differ phenomenally, but they are experiences of exactly the same atomic properties, and exactly the same complex colour and shape properties of individual objects. For that matter, they could

easily be experiences of the same particular objects as well. Yet the phenomenal difference between the two experiences is in some sense a matter of what they make phenomenally present to their subjects. To speak metaphorically, it is just a matter of the same 'pieces' being present but arranged differently in each case. In this and the following Section, I develop a new argument against property-based Hallucinatory Presence views, somewhat analogous to, but I think superior to the QIA. This argument is based on the idea that the structures made phenomenally present to us by our successful perceptual experiences are the sorts of structures we perceive in the world: ordinary physical objects standing to each other in spatial relations, rather than abstract objects such as Russellian propositions or complex properties.

In this Section, I begin the argument against property-based Hallucinatory Presence views, by laying out the exact sense in which our experiences make structured items phenomenally present to us. I argue that the phenomenal properties of our successful perceptual experiences depend on some of the items they make present to us being related in the right way, in particular, related so that they together form a particular structured item that is itself phenomenally present to us. From this it follows, I argue, that we can associate with each successful perceptual experience what I call a 'phenomenal binding condition'. These can be thought of as descriptions of what individual building blocks an experience makes phenomenally present to its subject, and how those 'blocks' are arranged into a particular structure. (I will give a more rigorous definition below).

I then argue that the possibility of cases like that of E_1 and E_2 where a pair of experiences differ in their phenomenal properties, solely because of differences in the structured items they make phenomenally present to their subjects, motivates a particular way of interpreting what it is for two experiences to differ qualitatively. On this interpretation, for two experiences to differ qualitatively, they don't have to make distinct objects, or atomic properties phenomenally present to their subjects. Rather, it is possible for two experiences to differ qualitatively *solely*

because they make different structures phenomenally present to their respective subjects, even if there is no difference in the properties and relations and/or individual objects they make so present. This is where I introduce the principle I call the '*extended* QRM'.

In the next Section, Section Three, I then use the results of this Section, to make a case that our successful perceptual experiences make phenomenally present to us fact-like structured items I call 'perceptible scenes' which have ordinary objects, bound together by spatial relations, as parts. I show that it follows from the extended QRM, that if a successful perception *E* makes such a perceptible scene phenomenally present to its subject, and a pure hallucination, *H* does not, *E* and *H* will differ qualitatively. Hence, it will jointly follow from the extended QRM and the claim that pure hallucinations do not make perceptible scenes phenomenally present to their subjects, that the HMT is false. From this it follows that property-based Hallucinatory Presence views are false, for the reasons I gave in the Introduction to this chapter.

I make the case for perceptible scenes being present in the successful perceptual case in the following way. Firstly, I argue that for any successful perceptual experience, *E*, there will be no gap between on the one hand, its phenomenal binding condition, *C* obtaining, and on the other hand, its accuracy condition, *A*'s obtaining, where an experience's accuracy condition is a proposition which specifies how the world would have to be for that experience to be veridical. I make my case for this claim in two ways. I make an intuitive case for this, based on the idea that a successful perception's phenomenal binding condition is just a full description of what that successful perception makes phenomenally present to its subject, whilst its accuracy condition *in the successful perceptual case* is really just a full description of the structure the subject of the experience perceives in the world. Those who endorse transparency should think there is no gap between the structures our experiences make phenomenally present to us, and the structures we perceive in the world. Secondly, I raise some further awkward questions for those views on

which the items phenomenally present to us in the successful perceptual case are not perceptible scenes, but rather structured items which have properties, but not objects themselves, as parts.

Before launching into the first stage of this argument, I will say a little about why I think it is not likely to be subject to the two problems with QIA that I diagnosed in the previous Section.

Obviously, the argument, if successful avoids problem 1). I explicitly argue in this Section for a principle which entails that, if I am right about what is present in the successful perceptual case, pure hallucinations cannot match successful perceptions in virtue of making properties alone (or structured items built up from properties but not objects), phenomenally present to their subjects. So unlike the QIA, my argument does not just assume a controversial claim about what similarities there must be between two experiences, with respect to the items they make phenomenally present to their subjects, in order for those two experiences to match. What about the second problem, that the QIA relied on claiming that something was just introspectively, or intuitively obvious, which was in fact clearly not obvious at all to many philosophers? Well, the major place that my argument relies on a claim about what is simply intuitive is when I defend the claim that a successful perception's phenomenal binding condition will entail its accuracy condition. It's certainly possible that what I say here won't strike all philosophers who read it as persuasive. But, unlike the claim that objects are phenomenally present in the successful perceptual case, this is not the sort of claim which we should have expected proponents of property-based Hallucinatory Presence views to have considered. So the mere fact that they don't discuss it as a potential problem for their view is hardly evidence that they would not find my claim about the relationship between accuracy conditions and phenomenal binding conditions intuitively compelling were they to consider it. After all, I'm going to spend a considerable amount of time setting things up before it's even clear exactly what a phenomenal binding condition is and why we should think experiences have them. Contrast this with the claim that objects are phenomenally present in the successful perceptual case, which is not only easily

grasped, but also clearly lurking in the literature. So at least we don't already have decent evidence that the key intuition which my argument relies on is not widely shared. Further, as I've already indicated, I do provide some extra support to the key claim that successful perceptions make perceptible scenes phenomenally present to their subjects, by showing that denying it has counter-intuitive consequences.

On then to the reasons for thinking that our experiences make structured items, not just pluralities of tropes, or atomic properties, or objects, phenomenally present to us.

In particular, I want to show the following: for a typical successful perceptual experience, *E*, its phenomenal properties partly depend on how some of the items it makes phenomenally present are related to each other, where the item *I* consisting of those items related in that way is itself phenomenally present to the subject of the experience. Further, *E*, or at least some perceptual experience, could have made a different structured item, constructed out of the same items phenomenally present to its subject. And had it done so, its phenomenal properties would have been different. For example, in the case of E_1 the 'building blocks' of the structure it makes phenomenally present to its subject, might include the red square, yellow rectangle, green circle and blue triangle themselves. Or[61] it might include the properties *redness*, *squareness, yellowness*, etc. Or it might even include not individual shape and colour properties, but rather the complex shape and colour properties instantiated by each of the perceived objects. Call the set of such items which E_1 makes present to its subject *Q*. The idea is as follows. Both E_1 and E_2 make all the members of *Q* phenomenally present to their respective subjects, and hence share some phenomenal properties. But they also make present to their subjects distinct structured, items, I_1 and I_2, and this is why they differ phenomenally. I_1 and I_2 are both constructed out of the same items, the members of *Q*. But although they both consist in the items of *Q*, related to each

[61] The or here is inclusive.

other in various ways, they are distinct items because the way in which the items of *Q* are related to each other, and hence bound into a structured item, differs in each case.

I don't think many of those philosophers who accept transparency would deny that structured items are made phenomenally present to us by successful perceptions. Nor, I think, would they deny that those structured items have as parts other items which are phenomenally present to us. Nor that in most cases another, distinct structured item, also capable of being made phenomenally present to us by a successful perception, could have been composed out of the same set of items. Certainly, everyone in the literature recognizes the importance of showing that their view is compatible with the obvious fact that experiences like E_1 and E_2 differ in their phenomenal properties. And every transparency-respecting view that I'm aware of accounts for this using the fact that E_1 and E_2 are related to different structured items, despite making the same (atomic) properties, and quite possibly the same objects, phenomenally present to their respective subjects. For example, as we saw, many proponents of property-based Hallucinatory Presence views, such as Tye and Dretske, are representationalists, and would hold that E_1 and E_2 bear different propositions as their representationalist contents. Similarly, amongst those who reject property-based Hallucinatory Presence views, Fish for example holds that *facts* are made phenomenally present to us in the successful perceptual case[62]. Obviously, there are separate facts associated with the world being as E_1 represents it to be and with it being as E_2 represent it to be, although if we take facts to be structured entities, as Fish seems to, those facts would have the same constituents. Property-based Hallucinatory Presence views which are not representationalist, but on which it is complex properties alone, and not objects, which are phenomenally present to us in the successful perceptual case, can also adequately distinguish between E_1 and E_2, I think, though probably at the price of treating the properties in question as

[62] (2009), pp.51-4.

properties of pluralities of objects[63], or regions of space-times, given that all the same complex properties of individual objects are present in each case.

I shall now produce an argument for the claim that they do indeed make structured items, which have other items phenomenally present to us as their constituent parts, phenomenally present to us[64].

Take the case of E_1 and E_2. Clearly they differ in their phenomenal properties. And it's hard to deny (at least if you accept the legitimacy of phenomenal presence talk at all) that this difference is a matter of there being a difference in what they make phenomenally present to their subjects. Now we've seen already that the difference cannot be a matter of which individual colour and shape properties or spatial relations the two experiences make present to their subjects. So what either E_1 and E_2 make present to their subjects cannot be a mere plurality of atomic properties and relations. What we want to show here of course, is that what distinguishes the two experiences is that they make present to their subjects different structured items, where those structured items have as parts other items the experiences make phenomenally present to their subjects. I will now argue for this, by eliminating other ways in which the phenomenal difference between E_1 and E_2 might be accounted for.

Plausibly the difference is in some sense a matter of what's going on with the items they make present to their respective subjects. If so any account of why E_1 and E_2 differ must instead cite facts about the items they make phenomenally present to their subjects. Given this, the following options seem to be exhaustive.

Firstly, technically speaking, E_1 and E_2 might not make atomic properties or individual physical objects phenomenally present to their subjects at all. Instead, they might only make present to their subjects single unstructured items that somehow encoded the objects and properties they

[63] Tye (forthcoming) takes this line.
[64] My argument bears some resemblance to Jackson's famous (1975) 'many properties' argument against adverbialism.

were perceptions of. For example, they might make facts or propositions, conceived of as unstructured items, phenomenally present to their subjects. But this seems to be ruled out by the fact that the subject of E_1 (or E_2) would be able to direct attention to some but not all of the individual items they make phenomenally present to their subjects. This suggests that what's present can't just be a proper partless item like an (unstructured) fact or proposition. For example, if you were undergoing E_1 you could certainly attend to the item(s) present to you connected to *redness squareness*, without attending to the items connected to *yellowness* and *rectangularity*[65].

So we know that E_1 and E_2 make *multiple* items phenomenally present to their subjects, but that they differ phenomenally and that their so differing is in some sense a matter of what's going on with the items they make phenomenally present to their subjects. How else might we account for the phenomenal difference between them, other than saying they make present to their subjects distinct structured items, I_1 and I_2, where each of I_1 and I_2 are themselves built out of other items which E_1 and E_2 respectively make phenomenally present to their subjects? Presumably there are two alternatives here. Firstly, in addition to the items which both E_1 and E_2 make present to their subjects, they might also make present to their subjects distinct items, I_1 and I_2 respectively, which are not structured items at all, but rather lack proper parts. Or secondly, the distinguishing items, I_1 and I_2 might be structured items, with proper parts; but none of their parts might themselves be present to the subjects of E_1 and E_2 respectively.

Suppose there are such distinguishing items, I_1 and I_2 connected to E_1 and E_2 respectively, which are either lack proper parts, or at least do not have other items which E_1 and E_2, respectively, make present to their subjects as parts. Now, either these distinguishing items are somehow

[65] Of course, attending to them will change the character of your experience, but there doesn't seem to be any reason in principle why you could not, in this case, be selecting amongst the items already phenomenally present to you, for some to receive more attention, and this is all we need to make the point that multiple items are so present.

connected to the shapes and colours and spatial relations that E_1 and E_2 are perceptions of, or they're not. I think trouble follows either way.

Suppose first they are not. My case against this is intuitive. Imagine you are undergoing E_1. Would there really be anything present to you separate from the colours and shapes (or colour and shape tropes, or coloured and shaped objects) phenomenally present to you, which distinguished E_1 from E_2[66]? I think it is pretty hard to say that there would be. The difference between E_1 and E_2 doesn't seem to be a matter of something they do in addition to making present to their subjects particular shape and colour properties (or objects of a particular shape and colour.) Rather it seems to have something to do with those shape and colour properties being 'arranged' differently in each case, even if we're not quite sure what the metaphor of arrangement comes to. Admittedly I do not have an independent argument to offer here. So this is one place where my argument is resistible. But we can add that the burden of proof is on anyone who takes this line to provide reasonable candidates for being the distinguishing items here.

The second option is to hold that, in fact, E_1 and E_2 do make distinct items phenomenally present to their subjects, but that those items actually are colour and shape and spatial properties or coloured and shaped, spatially related objects. This may seem hopeless; haven't we seen that E_1 and E_2 are experiences of exactly the same colours and shapes and spatial relations? But in fact, with enough ingenuity, we can find somewhat gerrymandered properties that distinguish them which are in some sense connected to the colours and shapes and spatial relations that their subjects perceive[67]. Suppose, for example, that the gap between the red square and the yellow rectangle in the case of E_1, is two feet. Then the red square perceived by the subject of E_1 will have at least one property that the red square perceived by the subject of E_2 does not, namely,

[66] I'm not here taking a stand on whether our successful perceptions make only colours, and shapes, spatial and motion properties phenomenally present to us, or whether instead, as Siegel (2010) thinks, they also make a variety of other properties so present, like causal properties, natural kind properties, and so forth.
[67] Thanks to Jeremy Goodman for discussion here.

being two feet from a yellow rectangle. On some views of properties, this will itself be a structured item, but on others it will not.

It seems doubtful, however, that such properties are phenomenally present to us in the successful perceptual case. All things being equal, we should not hold that highly gerrymandered properties are present to us in the successful perceptual case, without some independent reason to do so. This is especially plausible if we assume, as many proponents of property-based Hallucinatory Presence views do, that the properties our successful perceptions make phenomenally present to us are those they represent as instantiated in the world. It seems dubious that we would represent such gerrymandered properties, especially if we endorse a metasemantics for mental representationon on which it is easier for representational states to denote properties, the more natural those properties are in the Lewisian sense[68].

Further, if we conceive of properties like *being two feet from a yellow rectangle* as being unstructured, partless items, or as not making any of their parts phenomenally present to us, there is a second reason why E_1 and E_2 probably don't make such properties phenomenally present to their subjects. Suppose you were undergoing E_1 Clearly, you could, in principle, attend to the items phenomenally present to you connected to *yellowness* and *rectangularity* without being able to tell that any particular spatial relation was phenomenally present to you[69]. This strongly suggests that, even if *being two feet from a yellow rectangle* is phenomenally present to us, so is *being yellow and rectangular*. But now it looks like there's a somewhat suspicious double-counting going on. Why would successful perceptual experiences make present to us both a primitive spatial relation which codes in a colour and shape, and that colour and shape? Why not just make the colour and shape present to us, and some also less gerrymandered spatial

[68] See Sider (2011, pp.27-9) for an argument that this will be true on any reasonable metasemantics. Of course, colours are already very unnatural properties in the Lewisian sense. But presumably gerrymandered properties of being a certain distance from an object of a particular colour and shape, will be more unnatural still.

[69] This is shown by the fact that you could attend to to the items phenomenally present to you connected to *yellowness* and *rectangularity* without attending to anything that would tell you how far the yellow rectangle you were perceiving was from any other object

relation? If we assume, as proponents of representationalist property-based Hallucinatory Presence views do, that the properties which are phenomenally present to us in the successful perceptual case, are those which enter into the representational content of the successful perceptual experience, then I think the redundant representation here looks especially suspicious.

Again I don't think either of the arguments just given is irresistible. But together, I think they do have a fair amount of power.

So, we have eliminated the possibility that E_1 and E_2 make single, distinct items phenomenally present to their subjects. And we've eliminated the possibility that, although they both make multiple items present to their subjects, the items which distinguish them amongst the multiple items they each make present to their subjects are either proper part lacking, structureless items, or at least such that none of their parts are also phenomenally present to the subjects of E_1 and E_2 respectively. That seems to leave us with only the following option, given that the difference between E_1 and E_2 is a matter of what's going on with the items they make phenomenally present to their subjects: the items, I_1 and I_2, which distinguish E_1 and E_2, are both structured items, and items whose parts are themselves made present to the subject of E_1 and E_2. For example, E_1 might make present to its subject the existential Russellian proposition *that there is something red and square two feet to the left of something yellow and rectangular, which is two feet to the left of something blue and circular, which is two feet to the left of something green and triangular*. Or it might be a similar but *singular* structured proposition. Similarly it might be the following existential fact that E_1 makes present to its subject: **there is something red and square two centimetres to the left of something yellow and rectangular, which is two centimetres to the left of something blue and circular, which is two centimetres to the left of something green and triangular**. Or it might be a similar, singular fact into which the square, rectangle, circle and triangle themselves all enter.

Suppose we do treat facts or propositions as structured items that have objects and properties as parts, and hold that fact/propositions are present in the successful case, and so are the properties, and perhaps the objects which they have as parts. Then the problems outlined above disappear. Firstly, such views would be views on which E_1 and E_2 make multiple items phenomenally present to their subjects, so there's no problem with the fact that we can switch attention from one item phenomenally present to us to another. Secondly, the difference between the propositions or facts, which E_1 and E_2 respectively make phenomenally present to their subjects is partly a matter of how individual shape and colour properties or[70] shaped and coloured objects, are arranged into the relevant proposition or fact. So there's no worry that the phenomenal properties that differentiate E_1 and E_2 are somehow not properties they have partly in virtue of making colours and shapes (or coloured and shaped tropes, or coloured and shaped objects) present to their subjects. Those properties can be ones which E_1 and E_2 respectively possess in virtue of making those colours and shapes, arranged in a certain way into larger structures, phenomenally present to their subjects. And there doesn't seem to be any reason to think either that odd and gerrymandered properties have to be constituents of the facts, or that there must be any kind of 'double-counting' of the kind we saw above results from saying that a certain kind of gerrymandered complex properties are phenomenally present to the subjects of E_1 and E_2.

Note that I'm not claiming here that Russellian propositions and facts *exhaust* the possible candidates to be the structured items our experiences make present to us. For example, I haven't argued against the claim that complex properties of some sort or other might be the structured items phenomenally present to us in the successful perceptual case. More generally, there might be further candidates for being the sorts of structures that our experiences make present to their subjects beyond facts, Russellian propositions, or complex properties, for all I've said so far. What matters here is the following. By showing that our successful perceptual experiences make

[70] Inclusive.

present to us structured items, which have as parts items that are themselves phenomenally present to us, I have shown that successful perceptual experiences have what I call 'Phenomenal Binding Conditions':

Phenomenal Binding Conditions (PBCs): A successful perceptual experience *E*'s phenomenal binding condition *C* is a proposition which specifies the existence of all the atomic items[71] *E* makes phenomenally present to its subject *S*. Insofar as two or more of those atomic items jointly compose any further, structured items, which *E* also makes present to *S*, *C* specifies the relations which hold between those items (and perhaps others which are not themselves made phenomenally present by *E* to *S*) in virtue of which they compose that structured item. And *C* also specifies that those items are so related and compose such a structured item.

Intuitively, we can think of PBCs as complete descriptions of what's going on in a subject's phenomenal visual field, at a particular moment in time. Note though that I would not want to lean too hard on the metaphor, if metaphor it is, of the 'phenomenal visual field' in giving my official account of what a PBC is. To illustrate the idea with a simple example, suppose a subject, *S* has a successful perceptual experience *E* of a red square, which fills the whole of her visual field. And suppose, for example, the items phenomenally present to us in the successful perceptual case are singular Russellian propositions. Call the particular singular Russellian proposition which *E* makes phenomenally present to its subject *'R'*. The PBC associated with *E* will specify that there exists an object, *o* and the property, *redness*, and that *o* and *redness* are related to each other however they have to be related to each other[72] when they are combined in the way they must be to compose *R*. On the other hand, if we take what is present to us to just be the objects related as they are in the world, then the PBC associated with *E* will simply specify that there exists an object *o* and the property *redness*, and that they are bound together into a further

[71] Here an 'atomic' item, *I*, is an item such that *E* makes it phenomenally present to its subject without making any part of it so present. Note that this does not necessarily entail that *I* itself lack proper parts.

[72] And perhaps also further items which *E* does not make phenomenally present to its subject.

structured item by the relation of instantiation holding between *o* and *redness*. Of course, on any view the PBC associated with E_1 or E_2 will be much more elaborate than this.

Presumably, whenever two or more items compose a further item, there will always be some relation(s) they stand in to each other in virtue of which they do so (even if, at the most trivial end of things, these are just mereological relations like 'jointly composing'). If so, then there will be a non-trivial PBC associated with each successful perceptual experience; non-trivial, in the sense that it will specify how some structured items are composed, as well as which atomic items the successful perception makes phenomenally present to its subject.

Having shown that successful perceptions make structured items phenomenally present to us, we now need to give some account of what differences in the structured items two experiences make phenomenally present to their subjects make for phenomenal differences between those two experiences. For, at least if we take it that what E_1 and E_2 make present to their subjects are structured propositions, or facts, or fact-like entities constructed out of ordinary objects and their perceptible properties, and their constituents, E_1 and E_2 will make exactly the same properties phenomenally present to their subjects. Yet they clearly differ phenomenally.

Now, recall that in ch.1, I committed to only a single claim about when two experiences match, what I called the QRM. The QRM states that two perceptual experiences share all and only the same phenomenal properties, only if there is no qualitative difference in what they make phenomenally present to their subjects. I also said there that differences in which properties or relations[73] two successful perceptions make phenomenally present to their subjects count as qualitative. But I left the door open as to whether any other differences in the items two experiences make present to their respective subjects also count as qualitative. As it stands then, I said nothing that would explain why E_1 and E_2 differ phenomenally, given that there is no difference in the properties they make phenomenally present to their subjects. Here I will correct

[73] I also gave glosses involving tropes and objects, in ch.1, but we can suppress that complication here.

that, by giving an account of which differences in the structured items two experiences make phenomenally present to their subjects count as 'qualitative', and hence which differences lead to a phenomenal difference between the two experiences.

The best way, I think, to extend the idea of 'qualitative' differences to cover structural difference is as follows. We treat all differences between two structures as qualitative differences, unless they involve merely swapping one object in a particular 'slot' in a structure, for another, or one trope for another trope of the same property or relation. Perhaps swapping between properties, objects which bear those properties and tropes of those properties should also be held to make no phenomenal difference, but I will remain neutral on this. To see how this idea works, it will help to think of the structures phenomenally present to us in the successful perceptual case in terms of diagrams. These 'structure diagrams' represent at a high enough level of abstraction that the structures phenomenally present to us in the successful perceptual case can be represented by such diagrams whether they are Russellian propositions, or facts[74], or complex properties. These diagrams have both nodes, and lines connecting those nodes. The nodes are then labelled with predicates-a single node can get multiple predicates-whether for properties or relations. The lines between the nodes are labelled with relation terms only. If we think that only properties (or tropes) are made phenomenally present to us by our perceptual experiences the nodes represent as follows. Some nodes represent the properties which are present to us. Others may represent relations which are present to us but aren't actually amongst the relations which bind together the parts of the structure into that whole structure. For example, if the structures present to us in the successful perceptual case are existential Russellian propositions, then the relations which those propositions are partly about, are parts of those propositions, but they're not amongst the relations which bind together the parts of those relations into that proposition[75]. Suppose on the

[74] Or perhaps fact like entities which literally have objects related by spatial relations as parts

[75] For example, take the Russellian proposition *there exists a red square three feet to the left of a green circle*. Here, the relation *three feet to the left of* is part of the Russsellian proposition in question. But it's

other hand, we think that objects alone or objects and properties (or tropes) are phenomenally present to us. Then some nodes can represent the objects that are present to us, and the properties of those objects which we are present to us or which we experience them as possessing. Some of the nodes may instead represent relations which are part of the structure, but not themselves amongst the relations which bind together the parts into that structure. In order to indicate, what the nodes represent, we can add, at each node a *p* if the node indicates one or more properties, a *t* in the case of tropes, and an *o* in the case of particular objects, and if the node represents a particular relation, we can add an *r*. The labelled lines going between objects, meanwhile, represent the relations which bind together the parts of the structure the diagram represents into that whole structure.

Now, we can express what it is for two structures to differ qualitatively. The basic idea is that it's for the structure diagrams which can be used to correctly represent them to differ. If, however, we are uncomfortable with the idea that swapping between an object which instantiates a property, the property itself, and a trope of the property makes for a phenomenal difference, then we can hold that which 'letter' gets written beside a node is not a qualitative difference between structures. Similarly for swapping properties or relations for tropes of those properties and relations. As I've already said I don't see any good grounds for deciding whether such differences are qualitative, and so will remain neutral on this. Nothing in my argument below depends on this however.

Note that the sort of diagrams just described, even if we take the nodes on them to represent objects (or objects and associated properties (tropes)), don't actually carry any *de re* information about particular objects. Two structures which have entirely different objects amongst their parts could still be represented by the same diagram, provided that they were isomorphic to each other, since the nodes which represent objects (or objects and associated properties) don't

---

not amongst the relations which bind together the parts of that proposition into that proposition since Russellian propositions aren't literally *spatial* structures.

specify a particular object as the one filling the node. On the other hand, it will be obvious from what I said above that changes in which properties are part of a structure (or which tropes) does always make for a qualitative difference. I think that these two points are enough to show that my account of what differences between structures are qualitative is not totally unreasonable. Further, it's hard to see how else we could extend the notion. Below, I will therefore operate with this understanding of what it is for two structures to differ qualitatively, and will assume that when two experiences make present to their respective subjects, two structures, S_1 and S_2 which differ qualitatively, there must be a corresponding difference in the phenomenal properties of each experience. Anyone who wants to challenge me on this point needs an alternative account of when differences in associated structure make for differences in the phenomenal properties of two experiences. And they also need to show that on this account, the argument I give in Section Three below fails to go through. I think this is a tall order. The key claim I rely on below about which differences between structures make for phenomenal differences between the experiences they are associated with will, I think, come out true on any reasonable view of the matter.

Section Three-An Argument Against Property-Based Hallucinatory Presence Views

Here I will first argue that the items phenomenally present to us in the successful perceptual case are what I call perceptible scenes. Here a perceptible scene is an item, somewhat resembling a structured fact, constructed out of particular objects. Those objects are bound together into perceptible scenes by the actual sort of perceptible spatial relations we perceive in the world, and are also bound to properties by the instantiation relation[76]. I then show that the following follows

[76] Depending on one's view, perceptible scenes may *simply be* facts. I want to remain uncommitted on issues in the metaphysics of facts, however, so I will remain neutral on this.

from the extended version of the QRM argued for in the last section: if a successful perception, *E*, makes some perceptible scene, *K*, phenomenally present to its subject, then no pure hallucination *H* can match *E* simply by making some structured item constructed out of properties, but not objects, phenomenally present to its subject. Property-based Hallucinatory Presence views, are, by definition, views on which both the QRM and the HMT hold true, but pure hallucinations make no structures which have objects (as well as properties) as parts phenomenally present to their subjects. So if I'm successful in showing that perceptible scenes are phenomenally present in the successful perceptual case, and that the latter claim forms an inconsistent triad with the HMT and the QRM, I have refuted such views. In the next chapter, I consider the prospects for Hallucinatory Presence views on which pure hallucinations really can make perceptible scenes phenomenally present to their subjects.

Here then, is my argument in premise and conclusion form:

1) For any (visual) successful perceptual experience *E*, *E* has both a phenomenal binding condition *C* and singular and existential accuracy conditions A_S and A_E, and *C* entails both A_S and A_E.
2) IF, for a perceptual experience, *E*, the obtaining of its phenomenal binding condition entails the obtaining of its existential and singular accuracy conditions, THEN; *E* makes phenomenally present to its subject some perceptible scene, *K*.
3) For any (visual) successful perceptual experience *E*, *E* makes phenomenally present to its subject some perceptible scene *K* (from 1, 2).
4) If a successful perceptual experience *E*' of two or more distinct objects makes phenomenally present to its subject some particular perceptible scene, *K*, then no metaphysically possible pure hallucination *H* can share all the phenomenal properties of *E'* whilst possessing no further phenomenal properties.

5) For any (visual) successful perceptual experience *E*' of two or more distinct objects, there is no pure hallucination *H* which shares all *E*'s phenomenal properties and possess no further phenomenal properties. (from 3), 4)).

From 5), the falsity of the HMT obviously follows, since the HMT is a universal generalization, and hence supposed to hold for all successful perceptual experiences.

Here, accuracy conditions are conditions associated with experiences which specify the way a particular region of the world has to be in order for the experience to be veridical. For example, if I have a perceptual experience as of a red circle, whether successful perceptual or illusory or hallucinatory, then it will be accurate only if there is a red circle at a particular location in the world. I won't offer any independent defence of the view that perceptual experiences have accuracy conditions here. Firstly, this claim is already widely accepted and I find it highly plausible. Secondly, those who do seem to reject it, such as Travis (2003/2013) are not proponents of Property-based Hallucinatory Presence views, so my argument will have bite against actual advocates of such views. This is especially so, because the particular type of Hallucinatory Presence view endorsed by most advocates of property-based Hallucinatory Presence views is the Tye-Dretske view on which our experiences make phenomenally present to us the properties which enter into their representational content. Anyone who accepts that experiences have representational content must also accept that they have accuracy conditions. Finally, although I don't really think that any argument here is needed, I find Siegel's (2010, ch.2, especially, pp.30-33) argument for the claim that perceptual experiences have accuracy conditions extremely persuasive. I will assume that successful perceptions have both singular and existential accuracy conditions, although nothing really turns on this, since I think that if either the singular or the existential accuracy conditions of a successful perceptual experience *E* are entailed by the PBC of the experience, then it will follow that *E* makes present to its subject a perceptible scene.

Note that what premise 1) says is not that the claim '*E*'s phenomenal binding condition obtains' entails the claim '*E*'s existential and singular accuracy conditions obtain', but rather that for any successful perceptual experience, the particular proposition that is its PBC will entail both the particular proposition that is its existential accuracy condition, and the particular proposition that is its singular accuracy condition.

Now the inferences from 1) and 2) to 3), and 3) and 4) to 5) are clearly informally valid[77], so to show the argument is successful we only have to show that 1), 2) and 4) are true.

Premise 1) is supposed to be justified largely on the grounds that denying it is counter-intuitive. I shall attempt to bring out the relevant intuition below. I will give a quick argument for 2) below. And 4) is supposed to be supported by the conjunction of the QRM, and the claim that pure hallucinations make only properties, or structures which have properties but not objects as parts phenomenally present to their subjects. Since showing that 1) is true is the major task, I will briefly defend 2) and 4) and then concentrate on the defence of 1).

Firstly, 2), as a material conditional is true unless its antecedent is true and its consequent false. And it's hard to see how the antecedent of 2) could be true whilst its consequent was false. A perceptual experience's singular accuracy condition obtains when certain objects are a certain way, and are related to each other and the subject of an experience in a certain way. So a successful perception's PBC can only entail a singular accuracy conditions if it entails that the particular objects mentioned by the singular accuracy condition exist, and bear certain properties, and certain relations to each other and the subject of the experience. Given this, however, it's hard to see how we can avoid the conclusion that successful perceptions make perceptible scenes phenomenally present to their subjects. It's very hard to see how the PBC associated with an experience could entail that some perceptible objects exist, and instantiate certain properties and

[77] Since 3) makes a claim about all successful perceptions, it will be true of the particular subset of successful perceptions-those of two or more object-mentioned by 4).

stand in certain relations, unless it itself mentioned those objects, and specified that they bore certain properties and were related in certain ways. So a successful perception's PBC must, if it entails its singular accuracy condition, mention the objects mentioned by the singular accuracy condition, and specify that those very objects are a certain way. Now recall that an experience's PBC specifies how the atomic items a successful perception *E* makes phenomenally present to its subject are bound together into any structured item it makes present to their subjects. This means that saying that *E* has a PBC which specifies that some perceptible objects exist and instantiate certain perceptible properties and stand in certain perceptible relations is equivalent to saying that *E* makes present to its subject a structured item consisting of perceptible ordinary objects related by perceptible relations and instantiating perceptible properties[78]. Which is just to say that *E* will make present to its subject a perceptible scene. So 2)'s consequent follows from its antecedent, and hence 2) is true.

I think that, slightly rejigged, similar reasoning can be used to show that if an experience's PBC entails its *existential* accuracy condition, then it makes a perceptible scene phenomenally present to its subject. But since the antecedent of 2) says that an experience's PBC entails both its existential *and* its singular accuracy condition, I don't need the corresponding argument to go through with existential accuracy conditions in place of singular ones.

What about 4)? Well, recall that two structures differ qualitatively, on my view, if there is a difference in the relations which connect different nodes of their structure diagrams together. If our successful perceptual experiences make perceptible scenes phenomenally present to their subjects the lines which connect the points on their nodes will be spatial relations of some sort. For it is spatial relations which tie together objects into perceptible scenes. Spatial relations, however, cannot literally hold except between spatially located objects. So any structures which do not have objects as parts will differ qualitatively from those which are partly composed of

[78] Presumably the perceptible relations involved here will be spatial relations.

objects bound together by spatial relations. This means, that unless pure hallucinations themselves make structures which have objects as parts phenomenally present to their subjects, they cannot match successful perceptions which make perceptible scenes phenomenally present to their subjects. Or at least they cannot do so if we retain the QRM as a condition on matching. Note that this reasoning will not go through if the structured diagram associated with an experience has no connecting lines on it, as will happen in the case where I see only a single physical object, such as a red square, say. That is why 4) is restricted to perceptions of two or more objects.

So 4) goes through, given the assumption that it's impossible for pure hallucinations to make perceptible scenes phenomenally present to their subjects. I'll examine this assumption in the next chapter.

On then to the defence of 1). Obviously to establish 1), we must establish three things: firstly that all our successful perceptual experiences have (singular and existential) accuracy conditions, secondly that they all have PBCs, and thirdly, that for any successful perceptual experience, *E*, its PBC *C*, will entail both its existential accuracy condition A_E and its singular accuracy condition, A_s. Now, as I said above, I won't offer an independent defence of the claim that perceptual experiences have accuracy conditions here. And I've already defended the claim that they have PBCs. So that leaves us with just the last claim to establish.

Consider an arbitrary (visual) successful perception, say our old friend from the previous Section, E_1, a successful perceptual experience of a red square to the left of a yellow rectangle, to the left of a blue circle, to the left of a green triangle. Imagine that you are undergoing E_1. There is, of course, some particular set of items that E_1 makes phenomenally present to you, and those items are arranged into particular structured items. This can all be correctly or incorrectly characterized. As I said above, we can intuitively think of trying to characterize this as trying to get right what's happening in the 'phenomenal' visual field. Here the phenomenal visual field is defined as the

apparent area where the items phenomenally present to you are contained, rather than as the area of the world which is in your field of vision[79]. Now, just as you can correctly or incorrectly characterize how things are with respect to what items are phenomenally present to you and how they are arranged into structured items, so your experience can also characterize the world correctly or incorrectly. Furthermore, you can make at least a rough[80] description of how your experience characterizes the world to be. Now, ask yourself whether things could really be as they are with respect to what is phenomenally present to you-same items, arranged in the same way, things the same with respect to what is going on in the phenomenal visual field-without the world being the way that it would have to be for your experience to be accurate. It seems to me fairly obvious that the answer here is 'no'. Intuitively, there are not two distinct subject matters here which can be correctly or incorrectly characterized, one being how you perceive the objects in the world around you to be, and the other being how things are with respect to which items are phenomenally present to you, and how those items are arranged. Rather, these are just two different ways of picking out how the same group of atomic items are arranged into the same structure. If this is right, then clearly the obtaining of the experience's PBC, which says which atomic items the experience makes present to its subject, and how those items are related, will entail the obtaining of the experiences singular accuracy condition, which says how the items the experience counts as a perception of must be in order for the experience to count as accurate of the world. This also shows, incidentally, by the transitivity of entailment, that the PBC entails the experiences existential content too, since an experience's singular accuracy condition will always entail its existential accuracy condition.

I think that the intuition that there aren't really two different things here, firstly, the items which are phenomenally present to us and the structure they are arranged into, and secondly, the

[79] Nothing turns here on the legitimacy of talk of the 'phenomenal visual field' however, if it carries any further commitments beyond the idea that there are non-trivial phenomenal specification conditions.
[80] Perceptual content might be too fine-grained for you to be able to give an exact description of your experience's accuracy condition using the resources of any natural language. See Tye (2009, pp.103-9) for one defence of this claim.

structure we perceive in the world and the items which compose it, but rather simply one arrangement of items which can be picked out in two different ways, is hard to deny. It looks especially hard to deny if one is committed to the transparency thesis as of course proponents of Hallucinatory Presence views are. The intuitive idea behind the transparency thesis is that when we try to find some world-independent domain of items which are phenomenally present to us, we instead 'look through' the experience at the world around us. It's hard to see why this should be capturable by a formulation of transparency which just states a restriction on the types of objects or properties that our experiences make phenomenally present to us. After all a mere unstructured plurality of items isn't the sort of thing we perceive in the world. Nor is it clear that it helps much to say that the items present to us in the successful perceptual case are arranged into some structure which is *isomorphic* to the scene in the world which we perceive. Perhaps the best defence here is to claim that what is present to us is a complex property since complex properties are structures we perceive in the world. But note that, even on this view, there is still a gap between how the external world seems to the subject of an experience in virtue of their undergoing that experience, and how things seem when the subject turns her attention inward, and focuses on which items are phenomenally present to her, and how they are arranged. I think that this is the kind of gap which proponents of transparency ought to deny exists. It may be that we can explain away the intuition that there is no gap here. Those who think that experiences themselves and/or their properties are phenomenally present to us in the successful perceptual case will certainly think so. But if so, I think explaining it away will likely undermine the appeal of transparency itself, and will therefore be of little help to proponents of property-based Hallucinatory Presence views.

Note that it's important here that E_1 is a successful perception. Insofar as illusions and pure hallucinations also make anything present to their subjects, and so can have PBCs assigned to them, their PBCs cannot entail their accuracy conditions if they are not veridical. For the PBC associated with an experience is always a true proposition, whilst the accuracy conditions of

experiences that are not veridical are simply false. But this does not show that the intuition I've tried to bring out above is misleading *in the successful perceptual case*. Note also that the argument here doesn't depend in any way on us being able to tell, just by introspectively attending to them, what ontological category particular items phenomenally present to us fall under, whether they are objects or properties (or tropes), so it escapes problem 2) with the QIA, which we discussed in Section One.

In addition to the intuitive case just made for 1), I think we can also make a more direct case for 3). We can do so by drawing out some unpleasant consequences of various alternative views about what the structures our successful perceptual experiences make phenomenally present to us might be, other than perceptible scenes. Since the conclusion of the argument follows from 3) and 4) alone, even without 1) and 2), alternative sources of support for 3) are also support for the conclusion of the argument.

What are the alternatives to holding that perceptible scenes are phenomenally present to us in the successful perceptual case which the proponent of a property-based Hallucinatory Presence view might embrace? The most reasonable options I can see are the following, though I can't *prove* they're exhaustive. Firstly, they might hold that the structured items phenomenally present to us in the successful perceptual case are Russellian propositions (whether singular or existential)[81]. Or secondly, they might hold that they are complex properties, where the latter are conceived of as structured items built up from atomic properties[82]. Some of the problems I now discuss are specific to one or other of these views. Others I think, are likely to come up for any alternative to the view that perceptible scenes are the structures present to us in the successful perceptual case which is compatible with property-based Hallucinatory Presence views.

[81] Tye with his claim (2002) that an experience's 'phenomenal character' just is its representational content, seem to have, at one point, held something like this view. More recently, he has instead endorsed the view that complex properties are phenomenally present to us in the successful perceptual case (see forthcoming, p.16).

[82] Tye now endorses this view; see the previous footnote. Pautz (2007, p.498), attributes this view to several other philosophers.

A) Firstly, recall that we saw in Section One that Pautz (2007) claims that it's impossible for a (visual) perceptual experience to make any item *I* which is not spatially extended, phenomenally present to its subject. We saw that Pautz was probably overconfident in thinking that anything like this principle is *a priori*. Nonetheless it certainly seems as though we can appropriately characterize the items which are phenomenally present to us as standing in certain relations of distance and direction from each other. Consider, for example, two successful perceptions, E_3 and E_4, the first of a red square and a green triangle which look to be quite close together, the second of a red square and a green triangle which look to be rather further apart. Now, imagine undergoing, E_3 and then E_4 in short succession. And imagine that in each case, being a philosopher familiar with the jargon of current work on consciousness, you focused on what the experiences were like, phenomenally speaking 'on the inside', what they made phenomenally present to you. Wouldn't it be natural to say, that when you undergo E_4, the item(s) phenomenally present to you and connected to *redness* and *squareness* and *greenness* and *triangularity* respectively, are in some sense further apart within the visual field, than they are when you undergo E_3? Now the person who thinks that structures which have physical objects bound to each other by spatial relations as parts can take such talk of distance and direction within the phenomenal visual field literally. The items phenomenally present to you really are farther apart in one case than the other, and this shows up in what is phenomenally present to you, via the presence of a structured item that includes them bound together by an actual spatial relation. But as we've seen, proponents of property-based Hallucinatory Presence views typically hold that there are not literally any objects bound together by spatial relations phenomenally present to us in the successful perceptual case, since they typically hold that the structured items present to us in the successful perceptual case are a complex properties or existential Russellian propositions. Such proponents of property-based Hallucinatory Presence views therefore have to

somehow explain away the appropriateness of talk of 'distance' here. They must treat it as either mistaken-perhaps as a projection of the literally spatial relationships you would perceive between items in the external world onto the items phenomenally present to you-or as somehow metaphorical. But surely, all things being equal, the easiest way of explaining why the talk is appropriate is to treat it literally.

B) There is one very simple-minded objection to the claim that Russellian propositions could be the structured items phenomenally present to us in the successful perceptual case which I'm not sure how seriously to take. The basic idea behind the transparency thesis is that the only items which successful perceptual experiences make phenomenally present to us are the sort of items we can perceive in the world, such as ordinary objects and their properties. But Russellian propositions are not the sort of items we perceive in the world. So if we take transparency to imply that all *structures* phenomenally present to us in the successful perceptual case are structures we can perceive in the world, then it will conflict with the claim that Russellian propositions are phenomenally present in the successful perceptual case. This is another place where switching from focusing on the individual items present to us to the structures so present helps us see potential trouble for property-based Hallucinatory Presence views that we'd otherwise miss. How convincing this objection is depends on the degree to which this extension of the transparency thesis is supported by the same introspective data which supports the canonical formulation of the transparency thesis which I gave in ch.1

C) A second source of trouble for the view that Russellian propositions are phenomenally present to us in the successful perceptual case is this. We know that our successful perceptions make phenomenally present to us structured items which have other items phenomenally present to us as their constituent parts. I think a decent case can be made for saying that the relations which bind together those atomic parts into those structured items are themselves phenomenally present to us in the successful perceptual case. But

on the most prominent account of structured propositions, it's implausible that any of the relations which bind together the parts of structured propositions are phenomenally present in the successful perceptual case.

Why think that the relations which bind together the atomic items our experiences make phenomenally present to us are themselves phenomenally present to us? Well consider the point made under A) above that the items our experiences make phenomenally present to us seem to stand in relations of distance and direction within some sort of phenomenal visual field. Of course, in order not to beg the question against property-based Hallucinatory Presence views, we must leave open whether the relations here really are ones of distance and direction, and whether the apparent visual field here really is a region of space. But still, I think these relations of apparent distance and direction in the phenomenal visual field seem themselves to be present to us, whatever their true nature. Furthermore, the difference between, for example, E_1 and E_2, seems to be a matter of how the atomic items they make present to their subjects are distributed in the phenomenal visual field. This at least suggests that the holding between atomic items of the appropriate relations of 'distance and direction' here genuinely do play a constitutive role in binding together some of the items present to you into one of the structured items phenomenally present to you. And these relations of 'distance and direction' are, as I say, themselves phenomenally present to you.

Now consider the most well known account of Russellian propositions in the literature, that given in Soames (1987), on which they are simply sets. On this view Singular Russellian propositions are set theoretic entities built up out of sets that contain objects, properties and relations, or functions (such as the standard truth functions for the logical connectives) as parts. And existential Russellian propositions are built up from properties (and relations) and functions, but not particular objects. On Soames' way of constructing singular propositions, for example the singular Russellian proposition *David is mortal*,

consists of the set containing David and the set containing the property, *being mortal*: <David, <*being mortal*>>[83]. As this example shows there is often nesting of sets within sets on Soames' way of constructing such propositions. A simple existential Russellian proposition, meanwhile, such as the existentially quantified proposition *something is rich*, is identified with the set containing the property of being a non-empty set, and the function which takes each object and maps it to the singular proposition which says that that object is rich.

On this view, if such sets are the structured items our experiences make phenomenally present to us then strictly speaking the only relation which actually ties together the different elements of these sets is the set membership relation. But surely the set membership relation is not amongst the relations our successful perceptions make phenomenally present to us!

I think similar problems arise on some other accounts of structured propositions. For example King (2007) holds that they are certain kinds of *facts* involving relations between the things the proposition is about on the one hand, and the lexical items which make up a sentence expressing the proposition in question in some language on the other. I think that this view is also likely to lead to odd consequences if we think that the relations which tie together the parts of Russellian propositions are themselves phenomenally present to us in the successful perceptual case. But I lack space to argue for this in detail.

It's unclear to me whether the problem here applies only to the view on which the structured items phenomenally present to us are Russellian propositions or whether it also extends to views on which they are complex properties. No account of complex properties is as canonical as Soames (1987) account of propositions, making this hard to assess. Clearly, however, any set-theoretic account of complex properties, such as that

[83] Why the set containing the property being mortal, rather than just the property itself? Because, sometimes the same constituents can form multiple different propositions, and this is a way of coding in the distinction between such propositions without moving from sets to, for example, ordered n-tuples.

discussed by Lewis (1986, pp.31-3) will face the same problem here as Soames' (1987) account of structured propositions. Further we don't know for certain that the relations which bind their parts together on some *non* set-theoretic conception, will be the sort of relations we would be happy with holding can be present to us in the successful perceptual case. In addition, it's controversial whether there is any adequate non-set theoretic account of complex properties available in any case. Lewis (1986, pp.34-6) argues against the coherence of non-set theoretic accounts of structured complex properties on the grounds that the very same constituent properties would themselves have to appear as parts of the same complex properties multiple times. So it's at the very least far from clear that the view that structured complex properties are phenomenally present to us in the successful perceptual case, is out of the woods here.

I do not think that either the intuitive case I made for 1) above, or the three subsidiary arguments I made above in favour of 3), are successful beyond all doubt. The intuitive case made for 1) can certainly be doubted, A) is at most some, rather than very strong, evidence against views which deny 3). B) does not work against the complex property view. C) relies on the not completely obvious claim that the relations which bind together the parts of the structures our successful perceptual experiences make present to us, are themselves phenomenally present to us. And it *may* not work against the view that the structures phenomenally present to us in the successful perceptual case are complex properties. Still, I think that, taken as a whole, the argument of the last two Sections has built a decent case for the phenomenal presence of perceptible scenes in the successful perceptual case, and against property-based Hallucinatory Presence views. From now on, I will refer to the claim that our perceptual experiences make perceptible scenes phenomenally present to us as 'extended transparency'. I do so because it extends the original idea behind transparency, that what is phenomenally present to us in the successful perceptual case is the sort of items which we perceive in the world around us, to cover the structures which

are phenomenally present to us. In the next chapter, I'll discuss the prospects of theories on which perceptible scenes are made phenomenally present to us by pure hallucinations.

Section Four-Possible Objections to My Argument

Here I discuss three potential objections to the argument of the previous two sections. The first two, I think are not especially threatening. The last, however, *may* have some bite.

1) **Change Blindness Shows that Any Attempt to Imagine What Would Be Made Phenomenally Present to us Were we to Undergo a Particular Perceptual Experience Is Misleading**

A well known phenomenon in the psychological literature on perception is so-called change blindness. Change blindness occurs, roughly, when subjects fail to notice quite drastic changes in what they are perceiving in certain circumstances; see Simons (2000) for an overview of the literature on this. Now it is sometimes claimed (see for example O'Regan (2001)[84]) that what this shows is that the 'content' of our experiences is much more sparse than we typically think: we do not notice the changes, because in fact there is not a change in the content of our experience. Rather our experience only represents very coarse-grained properties which are preserved over even quite large changes in the picture we are looking at. Those who think that perceptual experiences have sparse content often claim that the reason we are surprised by change blindness[85] is that when we attend to any particular region in the visual field, the content of our experience about the region attended to becomes much less sparse. After attention, they say, it

[84] Also cited in this context by Tye (2009, ch.7).
[85] Subjects usually are surprised at the size of the changes that can occur without their noticing.

really does have roughly the level of determinacy with which we (wrongly) imagine our perceptual experiences represent everything within our visual fields, all of the time[86].

The relevance of this here is that, in the previous two sections, I relied in several places on an appeal to the reader to imagine what would be present to the subject of two imagined experiences, E_1 and E_2. But suppose we wrongly think that what is phenomenally present to us in the successful perceptual case[87] is much more determinate than it actually is. The worry is then that when we imagine what would be present to us when we undergo successful perceptual experiences such as E_1 and E_2 what we imagine is very different from what they would make present to us. If so, then it's hard to see how we could be justified in relying on such judgments about imagined cases.

Two points in response to this worry.

Firstly, the view on which change blindness actually shows that there is much less determinacy at the level of perceptual experience[88] is reasonably controversial. An alternative view[89] is that all the fact that we do not notice the changes shows is that, sometimes, items show up in the representational content of our experiences, or amongst the items phenomenally present to us, but we do not judge, at the level of thought that such an item has done so. The idea is that noticing the changes is a matter not just of perceptually representing the changes, or of there being a change in the items an experience makes phenomenally present to one, but also registering the change at the level of judgment about what one has seen and experienced.

Secondly, and much more definitively, recall the following. Proponents of the more radical interpretation of change blindness, on which it shows that our perceptual experience of the world

[86] O'Regan (2001).
[87] It is clear in O'Regan (2001), amongst other discussions of change blindness, that the surprising lack of determinacy is supposed to be a feature not just of the representational content of the relevant experiences but also of what they make phenomenally present to their subjects.
[88] Either in terms of representational content or in terms of what is phenomenally present to us
[89] See Dretske (2004).

is much less determinate than we typically think, usually explain the (purported) illusion of determinacy in the following way. When we perceptually attend to a particular region of the world, we really do they say come to represent that region with (roughly) the level of determinacy which we typically take our experiences of the whole visual field to have. Presumably it follows from this that if we imagine what it's like to successfully perceive some array of objects, whilst attending to them very closely, we are not (greatly) mistaken about the level of determinacy at which we would perceive them. Nor, presumably, will we be mistaken about the level of determinacy of what such an experience would make phenomenally present to us. This suggests, however, that, if we had simply stipulated above, that the coloured shapes the subjects of E_1 and E_2 experience take up a relatively small area of the visual field, and can all be attended to, simultaneously, then our judgments about what would then be present to the subjects of those experiences, would not be misleading, with respect to the aspects of what is phenomenally present to them connected to the relevant colours and shapes. We could then re-run the arguments of Section Two and Section Three, as arguments restricted to what goes on when we have successful perceptual experiences, and direct attention to a particular region of the visual field. This would mean that rather than having shown that all successful perceptual experiences make perceptible scenes phenomenally present to their subjects, we would instead have shown, technically, that a certain restricted class of experiences, those which involve attention to a particular region of the visual field, involve perceptible scenes being made present to us. But in fact, this restriction does not affect the argument against the HMT. The HMT is a universal generalization which is supposed to cover all perceptual experiences. So if it's impossible for pure hallucinations to match successful perceptions involving attention to some particular region of the visual field, then the HMT will come out false. Note also that, in any case, once we establish that perceptible scenes are sometimes present to us, it would be odd if, when our experience of a region of the world is less determinate, what was present to us was not a perceptible scene which had less determinate properties and relations as constituents, but rather

a proposition or complex property. So I do not think that the radical line on change blindness even gives us reason to think that perceiving something, but not attending to it, is a way of having a perceptual experience that makes something other than a perceptible scene phenomenally present to you[90].

2) **Phenomenal Binding Conditions are more Fine-Grained than Accuracy Conditions**

Famously, there can be differences in the phenomenal properties of two experiences, which involve representing the same object *o* as having one or other of two *necessarily co-extensive* properties. Peacocke (1992, pp.75-6)), in a slightly different context gives an argument which well illustrates this; there is a difference between seeing a regular diamond as a diamond, and seeing it as a tilted square. There may be a worry that this causes trouble for my claim that the PBCs of successful perceptions entail (both) their (singular and existential) accuracy conditions. On the one hand, the difference between seeing a shape as a diamond and seeing it as a square seems to be in some sense reflected in what is phenomenally present to us. If we accept transparency, it is hard to handle this without holding that different, though necessarily co-extensive, properties are phenomenally present to us in each case, since we cannot appeal to differences in the experience itself[91]. On the other hand, there might be some pressure to say that accuracy conditions, since they are meant to tell us in what circumstances an experience would be veridical, shouldn't be so fine-grained as to distinguish between necessarily co-extensive properties. But now the worry is that accuracy conditions will come apart from PBCs, because the former will be less fine-grained than the latter. The thought is that, there will be n-tuples of experiences with the same accuracy conditions, but different, although necessarily equivalent, PBCs, for each n-tuple of merely

[90] Though the lesson taught by change blindness is that attention to an item is a *necessary* condition on its being made phenomenally present to one, *this* would obviously entail that no perceptible scenes which are not attended to are made phenomenally present to us, even if they are consciously perceived.

[91] Maybe the proponent of transparency could avoid committing to a controversial hyper-intensional conception of properties here by appealing to a difference in the relational properties of the tilted square/diamond, in particular, a difference in the properties it has in virtue of being related to the subject of the experience in some way. It's not clear to what extent allowing this sort of manoeuvre would trivialize transparency.

hyperintensionally distinct properties that can be present to us. This is arguably in conflict with the idea, expressed in the last Section, that in some sense, there is no gap, in the successful perceptual case, between what structure an experience makes phenomenally present to its subject, and what structure it 'tells' the subject is located out there in the world.

We can distinguish two possible problems here. Firstly, there is the worry that if PBCs differ merely hyperintensionally, and accuracy conditions don't they will fail to entail them, and 1) will be false. The second is that if PBCs and accuracy conditions differ merely hyperintensionally, this somehow undermines the intuitive case that I made for 1).

On the first point, there is in fact no need to worry. 1) does not claim that, for each successful perceptual experience, its singular (or existential) accuracy condition is *identical* to its PBC. This would indeed be inconsistent with PBCs being more fine-grained than accuracy conditions. What it says instead is that for any successful perceptual experience, *E*, its PBC, *C*, will *entail* both its singular and its existential accuracy conditions, A_S and A_E. It's perfectly compatible with accuracy conditions being individuated in a less fine-grained manner than PBCs, that the former can entail the latter. So there's no direct counter-example to 1) here.

Secondly, we might worry that if PBCs are individuated more finely than accuracy conditions, this somehow undermines the intuitive case I make above for 1), by opening up a gap between on the one hand, what the constituents of the structures present to us in the successful perceptual case are, and how those constituents are arranged, and on the other hand which items have to be present and arranged a certain way in the world, in order for a successful perception to be accurate. Luckily, a good response to this worry is available. The fact that we can have experiences which differ phenomenally, depending on whether we see an object as a regular diamond or a tilted square is either consistent with transparency or it's not. If it's not consistent with transparency, we don't have to worry about whether it undermines my case for 1), since we already have an argument against property-based Hallucinatory Presence views, given that those

views *entail* transparency. On the other hand, suppose that it is consistent with transparency. Then it's hard to see how this could be so unless we have a hyperintensional conception of properties. Now suppose we have a hyper-intensional conception of properties. Then we will presumably also be able to define up structured propositions which can differ despite being necessarily extensionally equivalent and constructed out of properties and relations that are themselves necessarily extensionally equivalent. But then, even if treat *accuracy conditions* as being identical when ever necessarily extensionally equivalent, we could easily associate with each successful perception such a hyperintensional structured proposition. We could do this by deriving this proposition from the successful perception's accuracy condition, plus information about which merely hyperintensionally distinct properties the subject perceived various items to have. I think we could then run the intuitive argument for 1) as an intuitive argument that there is no gap between an experience's PBC and the accuracy condition derived proposition of this type associated with it.

3) **Objections Surrounding the Reification of Structures and Nominalism**

A basic objection from reification is this. I move very quickly from the conclusion that there is a phenomenal difference between E_1 and E_2 and that this difference is in some sense, a matter of how the atomic items they make phenomenally present to their subjects are related to each other, to the conclusion that there are structured items, I_1 and I_2 which consist of just those objects related in just those ways. But many will consider it ontologically extravagant to move straight from the claim that there are certain items related to each other in particular ways to the claim that there is an item which consists in those things, related to each other in those ways. I think that this objection has some bite, though of course it's only threatening if there is no way to re-write the argument of Sections Two and Three in a way that does not reify structures. It's far from obvious that the latter cannot be done.

For my own part, however, I am happy to accept the view that, whenever a plurality of items are related in particular ways and/or instantiate certain properties, there is some fact, or fact-like item consisting of just those items related in just those ways. The reason for this is that I endorse the following deflationary metametaphysical view: whenever a theory acknowledges that there exist some items standing in certain relations to each other, it is not a further cost to a theory that it acknowledges the existence of some structured item that consists in those items, so related.

A second objection to reifying the structures present to us in the successful perceptual case, starts from the fact that the nominalist view that there simply are no properties is a live option in metaphysics with many supporters. From this it follows that arguments are greatly weakened if their premises are in conflict with nominalism. However, some will feel that if there simply are no properties and relations, as nominalists claim, then it is hard to make sense of what is being claimed when it's said that perceptible scenes (or facts) literally exist as entities in their own right. After all the same plurality of objects are supposed to be able to compose more than one distinct perceptible scene (or fact) depending on, which properties and relations they jointly compose it with. It's arguably hard to make sense of this if there simply are no properties and relations, but only objects which can be one or another, and be related to each one way or another. This threatens my argument because it suggests that nominalism is inconsistent with the claim that there are perceptible scenes made phenomenally present to us by successful perceptions.

In fact, however, it's not very worrying if an argument against property-based Hallucinatory Presence views is inconsistent with nominalism. For those views are themselves inconsistent with nominalism. That is, if there are no properties, then pure hallucinations can scarcely match successful perceptions by making properties phenomenally present to their subjects. Instead, if they match successful perceptions by making items present to their subjects, presumably, those

items must be objects[92]. To be fair, strictly speaking, since property-based Hallucinatory Presence views only say that for each successful perception there is a *possible* pure hallucination which matches it phenomenally, they are inconsistent not with nominalism, but rather with the necessary truth of nominalism. But this qualification is unimportant, since it is usually assumed that if nominalism is true, then it is necessarily true[93]. In fact, I think that the apparent incompatibility of property-based Hallucinatory Presence views with nominalism is another serious mark against them, though I don't have space to go into this in detail here.

There is, however, a more subtle objection to my argument suggested by consideration of nominalism. Clearly, even nominalists must account for the apparently structured nature of what is phenomenally present to us in the successful perceptual case, and the way that experiences, such as E_1 and E_2 can differ phenomenally even though they are of exactly the same objects. And if they deny the existence of tropes, as well as of properties, they must be able to do so without bringing in tropes. Presumably the best way for them to do this is to introduce, in addition to talk of experiences making items phenomenally present to their subjects, a second, more complex notion of an item's being phenomenally present and experienced as being a certain way. This is in fact needed simply to cope with cases much more simple than E_1 and E_2. For example two different subjects can successfully perceive the same surface area to have distinct colour properties, one property more determinate, and the other less, if one subject has more acute colour vision than the other. The nominalist who also rejects tropes must explain the phenomenal difference between the resulting experiences as involving the same object being phenomenally present in each case, but experienced as being different ways.

[92] Or perhaps tropes, if they are trope nominalists. I think that the claim that pure hallucinations can make tropes phenomenally present to their subjects is likely subject to the same objections which I make to the idea that they can make objects so present.

[93] Though in fairness, I am not sure this should be assumed by those who accept some sort of link between conceivability and possibility. It is certainly not clear why, if there are no properties, we should be able to know this *a priori*. So perhaps there is still a slight worry here, on some views.

But now the following thought arises. Perhaps, even on a view on which nominalism is *false*, and there *are* properties, we could account for the phenomenal differences between E_1 and E_2 not in terms of differences in the structured items they make phenomenally present to their subjects, but instead in terms of a difference in the relations that they each respectively bear to the same unstructured plurality of items which they make phenomenally present to their subjects. If so, perhaps what is phenomenally present to us might be a mere unstructured plurality of individual items after all[94]. What appeared to be differences in the structures phenomenally present to the subjects of E_1 and E_2 could instead be differences in how each of those subjects is related to the members of that unstructured plurality. So long as those items present to us must be objects, this would still be no trouble for my argument, as property-based Hallucinatory Presence views would still be false. Suppose, however, that we allowed that the items in question could be properties. Perhaps, pure hallucinations could then match successful perceptions by making present to their subjects unstructured pluralities of properties, but also being related to those properties in the right, sometimes very complex, way. After all, the argument I gave above against property-based Hallucinatory Presence views assumes that what is present to us in the successful perceptual case is a structured item, rather than (in effect) putting the structure into how the experience is related to the items it makes phenomenally present to its subject. Once again perhaps the argument could be re-written in a way that avoids reification of structures. And if so, it might also block a property-based Hallucinatory Presence view of this sort. But it's certainly not immediately obvious how to re-write it.

Of course, to flesh out this idea into a serious proposal, we'd need to have some idea of what exactly it is to experience an item as being one way rather than another. And we'd also need to have some idea of how to interpret the QRM, given the assumption that what is present to us is an unstructured plurality of items after all, with experiences of the same such plurality differing

[94] On such a view, there might still *be* perceptible scenes. It's just that successful perceptions would not make perceptible scenes phenomenally present to their subjects.

phenomenally depending on how they are related to the items within it. Only when we know how to interpret the QRM in such a context, can we know just how things would have to be for a pure hallucination to match a successful perception, given that we are working in a framework where what is present to us is an unstructured plurality, and the structure is all in the relations an experience bears to it.

But there is both an obvious account available of what it is to experience an item as being one way or another, and a corresponding way of understanding the QRM on such an account. The obvious thing to do is to treat experiencing an item as being a certain way as a matter of *representing* that item to be that way. We can then adapt the idea that two experiences differ phenomenally when there is a difference in their associated structure diagrams in the following way. Firstly, the nodes now stand for the items which the experience makes phenomenally present to its subject and represents as being a certain way. Letters on the nodes, of *p*, *r*, *o*, or *t* represent whether the item at the node is a property, relation, object or trope. The predicates with which nodes are labelled now stand for ways that the items at the nodes are *represented* to be. And the lines and labels on them now represent ways that the items are *represented* as being related to each other. We can then as before, take experiences to differ phenomenally when they are assigned different structure diagrams, except in the case of differences in the letters at the nodes, about which we can again remain neutral. Since we're remaining neutral on this, pure hallucinations might be able to match successful perceptions even if the unstructured pluralities that the former make present to their subjects consist of properties, and the latter of objects. So we have a possible new type of Hallucinatory Presence view here. On this sort of picture, we can replace the claim that an experience's PBC entails its existential accuracy condition, with the claim that an experience's accuracy condition gives us what items it makes present to its subject and how those items are represented to be.

More generally, on *any* view of what it is for us to experience an item as being one way rather than another, we can represent the way a subject experiences an item to be, in the following way. A node in a structure diagram stands for the item, and the predicates assigned to that node, and the labels on the lines connected to that node to represent how it's experienced to be, and how it's experienced to be related to other items. We can then re-jig the key intuition of Section Three as the claim that for any successful perception, *E*, if all the items on the nodes on *E*'s structure diagram are as the diagram says they are experienced to be, then *E*'s singular (and therefore existential) accuracy condition will obtain. When so formulated, this doesn't imply that if an experience gets a particular structure diagram associated with it, the items which the nodes stand for, the items it makes phenomenally present to its subject, really instantiate the properties and relations which they are experienced as having. This last point, ought, I think, to reduce our confidence that the argument against Hallucinatory Presence views given in the previous Section can be successfully re-written to work against the type of view under consideration.

I think that this problem for my argument has a decent amount of bite. However, I do have a point to make in response which I think lessens the attractiveness of views which put the structure of perceptual experience into the relationship between the experience and an unstructured plurality of items. Before proceeding to that response, I also want to note that, even if this objection to my argument succeeds, the type of property-based Hallucinatory Presence view it results in will differ greatly from those extant in the literature. Those extant in the literature, or at least those which are worked out in any detail, seem to presume that when an experience is phenomenally conscious, the phenomenal properties which that experience has are fixed by its getting some structured proposition and complex property associated with it. None of them seem to hold that what is present to us is an unstructured plurality of items, with the phenomenal properties of an experience depending on *how* it is related to that unstructured plurality. This means that even if this particular objection to my argument succeeds, the upshot of that argument is still interesting. For it will have shown that we should favour a new and as yet

under-explored type of property-based Hallucinatory Presence view, over more established views of that type. The fact that the view is under-explored of course, ought to *somewhat* undermine confidence in it, and hence in the objection to my argument currently under discussion. Nasty surprises may lurk for views which have not yet been picked over in the literature.

My basic response to the objection under consideration here is this. The proposed view is that pure hallucinations can match successful perceptions, because the pure hallucinations make present to their subjects unstructured pluralities of properties (and relations), but the phenomenal properties of perceptual experiences depend partly on *how* they are related to the items they make phenomenally present to their subjects. Here, I propose to argue against this view by the following steps. Firstly, I will give some brief reasons for thinking that the relevant relation between experiences and items that's doing the work here is a representational one. Experiencing an item to be certain ways is a matter of representing it to be those ways. I shall also argue that in the successful perceptual case, the items which we represent to be certain ways, and which hence are phenomenally present to us are objects, and not properties themselves. This tells us that even if what is phenomenally present in the purely hallucinatory case is an unstructured plurality of properties, what is phenomenally present to us in the successful perceptual case will be an unstructured plurality of objects. From this, I argue that it follows, given the QRM, that even if pure hallucinations can match successful perceptions, the particular pure hallucinations which match successful perceptions will not at all be those we would pre-theoretically expect to. Meanwhile those pure hallucinations which proponents of the HMT would antecedently expect to match successful perceptions will not. I argue that this severely limits the amount of support the view under discussion gets from the fact that it preserves the HMT. This shows that it may not do any better than Anti-Presence or Phenomenal Difference views which allow that successful perceptions make whole perceptible scenes phenomenally present to their subjects.

Why think that to experience an item a certain way is to represent it as being that way? Firstly, if experiencing an item to be one way or another is something we can do even in the purely hallucinatory case where we perceive nothing, it's hard to see what else it could be, other than representing that item to be one way or another.

Secondly, we should accept that in the successful perceptual case, the unstructured plurality of items which are phenomenally present to us, and which we represent to be one way or another, will consists of objects. For the view that our successful perceptions represent particular objects to be certain ways is widely accepted, and it seems *ad hoc* to depart from it just to save the particular type of property-based Hallucinatory Presence view we are considering. Furthermore, recall that we still want to be able to say that for any successful perception *E*, and structure-diagram correctly assigned to it, *D* that if the items at *D*'s nodes have the properties it assigns to them (and are related to each other by the relations the lines on *D* are labelled with), then *E*'s (existential and singular) accuracy conditions will come out true. These accuracy conditions, whether or not they are the same thing as *E*'s representational content are also usually thought to entail that there are objects in the world with certain properties.

Now, recall, that since the view we are considering is supposed to be a property-based Hallucinatory Presence view, it will be a view on which the QRM comes out as true. This means that for a possible pure hallucination, *H* and a successful perception *E* to match, there must be no *qualitative* difference in the structure diagrams, D_H and D_E that are assigned to each. Now, since *E*, being a successful perception, represents particular objects as being certain ways, the predicates which the nodes on D_E get labelled with will be the sorts of properties we perceive ordinary objects to have, such as, at a minimum colour and shape properties. And the relation-terms which go on the lines on D_E will be the sort of relations we perceive to hold between ordinary objects, such as spatial relations. But now, for D_E and D_H to match qualitatively, they must have the same property terms labelling their nodes, and the same relations terms labelling

the lines between nodes. So the predicates and relations terms on D_H will be terms for the sorts of properties and relations we perceive in the world, such as colours and shapes, and spatial relations. But at the same time, we know that predicates and relation terms on D_H stand for the properties and relations which *H* represents the items it makes phenomenally present to subject to have. And we know that the items *H* makes present to its subjects are properties. So *H* must represent *properties* as possessing the sorts of properties which we perceive *objects* to instantiate in the successful perceptual case. And it must represent *properties* as related to each other in the ways in which we perceive *objects* to be related to each other through our successful perceptions.

However, I think this conclusion is very damaging for the type of property-based Hallucinatory Presence view being discussed here. The picture of things we now have is one where pure hallucinations can match successful perceptions only by representing properties themselves to be the ways, such as being coloured and shaped and spatially related, which successful perceptions represent objects to be. From this it follows that pure hallucinations which represent that there are some objects which are a certain way, cannot match successful perceptions. But this picture of things is far enough away, I think, from the basic intuition that supports the HMT, that it's not clear that it gains any support from that intuition. We saw near the beginning of ch.1, that it's unclear how best to formulate that intuition, but that the HMT itself is probably in fact too weak. And certainly the relation-based view sketched here will be inconsistent with the stronger formulation of that intuition in counter-factual terms which I suggested might best capture the intuition behind the HMT. If to match a successful perception, a pure hallucination must represent *properties* as possessing perceptible properties and/or being related by perceptible relations, it will not be true that of every (or most) actual successful perceptions that had they been turned into pure hallucinations by destroying the objects of perception and preserving their neural base, their phenomenal properties would not have changed. After all they would hardly come to represent that *properties* were coloured or shaped, or spatially related, in such

circumstances. But the weirdness here goes beyond just the denial of the counter-factual version of the HMT. It also involves the denial that hallucinating an object is ever a way for pure hallucination to match a successful perception. This means that on this view, necessarily: those hallucinations which share the existential content of successful perceptions cannot match them phenomenally. Instead, they can be matched phenomenally only by pure hallucinations of a certain odd type with very strange representational content. But surely the basic intuition behind the HMT involves the idea that it is in the cases where the world perceptually seems the same to the subject of a successful perception and a pure hallucination, that matching occurs. And surely we also think, at least if we are in the grip of the intuition behind the HMT, that if how things perceptually seem differs greatly for the subject of a successful perception, *E* and a pure hallucination, *H*, that *H* and *E* will not match.

This, I think, suggests that what the view gets us is the HMT only, in effect on a technicality. For the particular hallucinations which, intuitively, we would think would match particular successful perceptions will not. If so I think it does not really gain an advantage over the view that perceptible scenes are present in the successful perceptual case from its ability to preserve the HMT.

Chapter Three: Can We Hallucinate Perceptible Scenes?

Introduction-

In the previous chapter, we saw that if it's (metaphysically) impossible for pure hallucinations to make perceptible scenes phenomenally present to their subjects, then all Hallucinatory Presence views are false, and we will be forced to choose between Anti-Presence and Phenomenal Difference Views. Many readers will, I suspect, think that this is indeed impossible. For many philosophers seem to hold that it's impossible for pure hallucinations to make particular objects, or items which have particular objects as parts (as perceptible scenes do) phenomenally present to their subjects. Here I examine this assumption. I argue that, although something like this assumption seems to be tacitly assumed in a lot of the literature, it's very unclear what the justification for it actually is. I show that most of the quick and easy arguments for it in fact fail. It turns out to be much trickier to refute the view that ordinary physical objects, or items, such as perceptible scenes, which have ordinary physical objects as parts, can be present in the purely hallucinatory case than is commonly thought. Eventually, however, I will give my own argument against the view that this is possible which does I think show that any such view will have reasonably severe costs. Given the considerable difficulties which beset all of the obvious rival views extant in the literature I don't think my argument shows beyond doubt that such views are false. However, I think it does raise enough worries for such views to make it worth examining

whether Anti-Presence and Phenomenal Difference views have, overall, lower costs than the radical Hallucinatory Presence view on which perceptible scenes which have particular physical objects as parts, can be present in the purely hallucinatory case. It is to such Anti-Presence and Phenomenal Difference views that I turn in the next chapter.

My plan of attack for this chapter is as follows. In Section One I give evidence that the literature has at least tacitly assumed that it is impossible for hallucinations to make particular objects, or perceptible scenes, phenomenally present to their subjects. Then, in Section Two, I sketch a picture of what phenomenal consciousness and phenomenal presence are, on which it *is* possible for pure hallucinations to make ordinary physical objects, and perceptible scenes which have ordinary physical objects as parts, phenomenally present to their subjects. I call this view the 'Quasi-Descriptivist' view of pure hallucinations, because it *somewhat* resembles descriptivism about proper names. I argue that, if some of the key conclusions of ch.1 and ch.2, stand up, then this view has certain attractions which ought to at least make us take it seriously. In Section Three, I deal with a series of obvious objections to the view, arguing that in fact, none of these objections is successful. In Section Four, however, I give my own argument against the Quasi-Descriptivist view. This argument is also intended to work against other views on which it is metaphysically possible for pure hallucinations to make perceptible scenes which contain ordinary objects as parts phenomenally present to their subjects. It thus completes my argument against Hallucinatory Presence views. For we saw in Section Three of the last chapter that I need only the assumption that pure hallucinations cannot make perceptible scenes phenomenally present to their subjects to complete my argument against such views.

Section One-A Neglected Option in the Philosophy of Perception

We saw in Ch.1 that there is a tension between transparency and the HMT. This tension arose as follows. The HMT seems to require that for each successful perceptual experience, *E*, there is a pure hallucination *H*, such that what *H* makes phenomenally present to its subject does not differ qualitatively from what *E* makes phenomenally present to its subject. But transparency threatens to show that ordinary physical objects are present in the successful perceptual case and it's not clear how pure hallucinations could come to make such objects present to their subjects. In the previous chapter we saw how to turn this tension into an argument from transparency against the HMT, based on the extended transparency thesis, the claim that our successful perceptions make perceptible scenes (which have ordinary physical objects as parts) phenomenally present to us. Now, the problems raised by transparency and extended transparency for the HMT could easily be resolved if we simply held that pure hallucinations can make present to their subjects ordinary physical objects, and items, such as perceptible scenes, which contain ordinary objects as parts. However no one, to my knowledge, has explicitly defended a view of this sort[95] . I think the tension between the HMT and transparency is now widely recognized. If so the lack of defenders of this sort of view is strong evidence that philosophers of perception have at least tacitly assumed that pure hallucinations cannot make particular objects (or, presumably, items which have them as parts) phenomenally present to their subjects. The result of this assumption, has, I think, been a string of implausible views which are caught between two stools. On the one hand philosophers have tried to be true enough to the introspective evidence which motivates transparency to not be refuted by it. Yet they have also tried to hold that the items which are

[95] In the case of pure hallucinations. Alston (1999, pp.191-2) defends the view that, in the case of impure hallucinations, where one successfully perceives some of the spatial region around oneself, but hallucinates one or more objects within it, empty space is phenomenally present to one, but 'appears' to be the hallucinated object. But in the case of *pure* hallucinations, Alston's view is that they make phenomenally present to their subjects some sort of mind-dependent mental image. Lewis (1983) may seem to suggest a view on which hallucinations (pure or impure) can make ordinary physical objects in other possible worlds phenomenally present to their subjects, which, given Lewis' modal realism, would imply that they make ordinary mind-independent objects phenomenally present to their subjects. I think a careful reading of Lewis, however, indicates that all that he is committed to is the view that, when someone has a hallucination as of a particular object, *O*, one sort of representational content possessed by their experiences can be given by the set of possible worlds where they are perceiving a counter-part of *O*.

present in the successful perceptual case are all of a type such that it's possible for pure hallucinations to make items of that type phenomenally present to their subjects. Hence, in the first third of the 20th century, sense-datum theorists posited that there were items present in the successful perceptual case, which were not ordinary physical objects, but which shared with ordinary physical objects the ability to instantiate colour and shape properties. And, as we saw in the last chapter, more recently, proponents of property-based Hallucinatory Presence views have defended the view that bare properties can be phenomenally present to us in perception without any bearers of properties being so present[96]. A.D. Smith (2002) has even defended the view that non-existent objects are phenomenally present to us when we hallucinate (purely or impurely). If we can work out a view on which it's possible for pure hallucinations to make perceptible scenes phenomenally present to their subjects, then we can avoid all of these unpleasant views, without being forced to deny transparency or extended transparency, or the HMT, or the QRM. In the next Section, I will sketch one such view, and give some reasons for thinking it might be true.

Section Two-The Quasi-Descriptivist View of Hallucination

We saw in ch.2 that we can't simultaneously hold onto the QRM, the HMT and extended transparency. From this the following thought naturally arises. If there is good evidence for each of these claims, but the only way to render them consistent is to hold that some possible pure hallucinations make perceptible scenes phenomenally present to their subjects, then perhaps we should simply endorse the view that pure hallucinations can make perceptible scenes phenomenally present to their subjects! This would involve giving up on the idea that it's impossible for pure hallucinations to make items which have objects as parts phenomenally

[96] Though we saw in Section One of ch.2 that even if objects are present in the successful perceptual case, and cannot be present in the purely hallucinatory case, this does not necessarily refute the HMT, absent a further controversial assumption that this difference entails a phenomenal difference.

present to their subjects. But before we reject this strategy for that reason, we should surely examine just how strong our reasons for thinking that this is impossible are. At the very least, evidence for the combination of extended transparency, the HMT, and the QRM, is at least prima facie evidence for the view that it's possible for pure hallucinations to make perceptible scenes present to their subjects. In this Section, I explore one such view on which pure hallucinations can indeed do this, and show that it has some attractive features.

The availability of such views can I think, be blocked by the following line of thought. Pure hallucinations are usually thought of as experiences which are not causally hooked up to the external world with its ordinary objects and perceptible scenes, but rather as purely internal goings on. Now, strictly speaking, this is not quite right of course, since there can be pure hallucinations which are caused by external objects, even by external objects which are constituents of a perceptible scene which fits the existential accuracy conditions of the hallucination, provided the causation is deviant in some way. Still, the image of a pure hallucination as being unmoored from any objects or scene in the world is the standard one. Guided by this image it becomes easy to think that for pure hallucinations to make particular objects phenomenally present to their subjects would require something magical and arbitrary to happen. The worry is that if some possible pure hallucination, *H*, makes phenomenally present to its subject a particular object *O*, or a particular perceptible scene, *S*, there will be no underlying facts in virtue of which the fact that: *H makes O (or makes S) phenomenally present to its subject*, obtains. If however, we do not want to treat phenomenal presence facts as somehow primitive, then we will not want our phenomenal presence talk to pick out a relation *R* such that at the actual world it can be a brute primitive fact that *R* holds between an experience and some object or scene.

This worry about arbitrariness I suspect lies behind the feeling that it's just *obvious* that pure hallucinations cannot make particular objects or perceptible scenes phenomenally present to

their subjects. Now, there are a couple of ways in which I could question this line of thought without denying that if a pure hallucination*H* is not caused by any physical external object, *o*, then there is no non-arbitrary way to associate any such particular perceptible scene with *H*. Firstly, someone might suggest the following. In those unusual cases where objects in the world cause pure hallucinations in some deviant way, those objects, or at least perceptible scenes which have them as parts, really are phenomenally present to the subjects of the experience. And this is enough to make the HMT come out true. But *typically* purely hallucinatory experiences do not make particular objects or perceptible scenes phenomenally present to their subjects. Secondly, perhaps things go radically vague here[97]. In particular, one might hold that, for each pure hallucination, it is determinately true that it makes some perceptible scene, and some certain number of objects phenomenally present to its subject, but for every perceptible scene it is vague whether or not that hallucination makes it phenomenally present to its subject.

But in fact, we don't need to take either of these lines. For there is an alternative account of the underlying facts in virtue of which pure hallucinations make particular scenes phenomenally present to their subjects, which does not rely on any causal connection holding between that scene or some or all of its constituents, and the pure hallucination in question. I call this account the 'Quasi-Descriptivist view of (pure) hallucinations'. Many philosophers think that perceptual experiences, including pure hallucinations, carry existential representational content. On the quasi-decriptivist view we simply take the existential content of an experience to determine which items, and in particular which perceptible scene and which objects, it makes phenomenally present to its subject, in roughly the way that descriptivists think that, for each name in the language, an associated definite description determines the reference of that name[98]. The idea is that each experience bears some proposition as its existential content. This proposition would take the form of specifying that there existed some object or objects with certain properties

[97] Thanks to John Hawthorne for suggesting this.
[98] Technically of course, at least since Searle (1958), many descriptivist have believed that it is a 'cluster' of descriptions, rather than a single description that determines reference.

related to each other and to the subject of the experience in certain ways. Unless the experience bears impossible content this proposition will clearly determine a possible perceptible scene, since perceptible scenes are just items consisting of ordinary objects instantiating certain properties, and being related to each other in certain ways. Or rather, strictly speaking, the content will determine a class of possible perceptible scenes which are all qualitatively identical to each other. *If a scene within this class exists at a world, and only that scene does, then we can treat the experience as making present to its subject that particular scene*. Since our experiences make present to us not merely structured items, but also parts of those items, on the Quasi-Descriptivist view, when an experience makes a perceptible scene present to its subject, it also makes the objects and properties and relations which are its constituents so present. If multiple qualitatively identical such scenes exist, things become a bit trickier, but there are options. Maybe in such cases it's determinately true that the experience makes a scene of the relevant type phenomenally present to its subject, but for each scene *S* of the relevant type, it is vague whether or not *S* is phenomenally present to the subject of experience. Or maybe we use something like physical closeness and direction from the perceiver to break the tie, choosing the scene which is closest to the subject of the experience, using direction to break ties in distance. This latter option would fit better than it might seem to with the general idea behind the Quasi-Descriptivist view, which is that the representational content of an experience determines which items it makes phenomenally present to its subject. For presumably, assuming they carry representational content at all, even pure hallucinations carry singular content about their subjects, and also have existential content which locates the items perceived relative to the subject of the experience.

What about the case where there is a pure hallucination *H* with existential content, *C*, such that there is no perceptible scene which exactly fits *C*? Here it gets a little trickier to decide what the proponent of the Quasi-Descriptivist view should say *H* makes phenomenally present to its subject. Firstly, we might introduce a comparative notion of being more or less like the members

of the class of qualitatively identical scenes which do perfectly fit that existential content. We could then declare that *H* makes present to its subject the scene which is most like the members of the class of scenes which fit *C* perfectly. After all, some perceptible scenes in the world really do seem more alike than others, and in an 'all-things considered' sense of more alike. For example, two mountain vistas involving mountains of similar shape, both at sunset, will be quite similar to each other, as these things go. And they'll be much more similar to each other than to a view across a flat part of the Sahara. Of course, some will be suspicious of such a notion unless it can be precisely defined. But it's hard to assess how much bite this sort of complaint has here, and I would not wish to rest my opposition to Hallucinatory Presence views on it.

Of course, it's hardly likely that there is a single, perfectly precise metric of similarity between perceptible scenes which is the 'correct' one here, and which therefore really determines exactly which scene such a pure hallucination makes present to its subject. Rather, presumably there is either one fairly vague notion of all things considered similarity here, or it's vague which of several different precise notions is the one which does the work of determining which scene a pure hallucination makes phenomenally present to its subject, or both. Either way, it will often come out as vague *which* perceptible scene a pure hallucination makes phenomenally present to its subject. Even when this goes vague though it may be perfectly determinate that there *is* a scene that it makes present to its subject. And it may be determinately true as well that there is a scene which is phenomenally present to the subject and has properties which fall within some particular range on various dimensions. But it's not clear that some vagueness here is a problem, absent some argument for the claim that phenomenal presence facts are always determinate.

Alternatively, if we think that we don't have a grip on any kind of all-things considered metric of similarity between perceptible scenes, or if for some reason we don't like vagueness here, we might instead simply say that if no perceptible scene at a world perfectly fits the existential content of a metaphysically possible pure hallucination *H*, then *H* simply does not make any

perceptible scene phenomenally present to its subject. On such a view, it is metaphysically possible for pure hallucinations to make particular objects phenomenally present to their subjects, but this requires that the world co-operate in quite a specific way, by providing a perceptible scene whose existence is sufficient to guarantee the truth of the pure hallucination's representational content.

Whatever option we take in cases where there are either no, or multiple scenes which exactly fit the existential content of a pure hallucination, however, we can now specify the general shape of the Quasi-Descriptivist view. The basic idea is simply that the existential content of an experience can determine which perceptible scene a pure hallucination makes phenomenally present to its subject. This gives us an alternative non-arbitrary way in which the underlying facts about a pure hallucination can determine which perceptible scene it makes present to its subject, when the sort of causal connection that there is between the constituents of a particular scene and the experience in the case of successful perception, is absent. On all versions of the quasi-descriptive view, whenever there is a unique perceptible scene at a world which perfectly fits the existential content of a pure hallucination, that pure hallucination will make that perceptible scene phenomenally present to its subject[99]. It's merely in cases where there is no unique matching scene that the various ways of working out the details of the view discussed above, differ. So, if the Quasi-Descriptivist view is true, in every metaphysically possible world where there is a pure hallucination, *H* and there is some unique perceptible scene *S* which exactly fits *H*'s existential content, *H* will make *S* phenomenally present to its subject. If we add the plausible assumption that for any successful perception, *E*, there is some metaphysically possible world, *w*, where there is both a pure hallucination, *H* which shares *E*'s existential content, *and* a single unique scene which fits the existential content of *H* exactly, then it's easy to see why the Quasi-Descriptivist view allows us to preserve the HMT, the QRM, and extended transparency. Since *E* is a successful

[99] And presumably also some of its parts, such as some of the objects within it, though for space reasons I won't try to sketch exactly how this might work.

perception, the scene S_1 that it makes present to its subject, will fit that existential content exactly. Now, given the Quasi-Descriptivist view, *H* will also make present to its subject a scene S_2 which perfectly fits *H*'s existential content. Since *E* and *H* share exactly the same existential content, however, S_1 and S_2 cannot differ qualitatively; no two perceptible scenes can differ qualitatively, and yet fit exactly the same content[100]. If there is no qualitative difference in the perceptible scenes *E* and *H* make phenomenally present to their subjects, however, then QRM is perfectly consistent with their matching phenomenally. And since there is such a metaphysically possible *H* for any *E*, the QRM will also be consistent with there being for *every* successful perception, a hallucination which matches it, or in other words, with the HMT itself. And all of this is obviously consistent with the claim that successful perceptions always make perceptible scenes which have ordinary objects as parts phenomenally present to their subjects. So the Quasi-Descriptivist view allows us to reconcile the HMT, the QRM and extended transparency.

Above, I've already said that, since the QRM, extended transparency, and the HMT are independently plausible, this gives us some evidence in favour of views which enable us to reconcile them. Since the Quasi-Descriptivist view allows us to reconcile them, and it's not clear how else to do so, this gives us some evidence in favour of the Quasi-Descriptivist view.

The case for the Quasi-Descriptivist view is strengthened when we see that it's far from being a view on which there is no connection between what determines which items an experience makes phenomenally present to its subject in the successful perceptual case, and what determines what items it makes phenomenally present to its subject in the purely hallucinatory case. Instead, if we squint a bit, we can see it as a view on which what goes on in both cases is actually the result of the workings of a single principle. This is reasonably important. Suppose the

[100] Any difference in their structure diagrams which counts as qualitative will be a difference either in which properties are associated with a node representing an object with the scene, or which relations get assigned to the lines which represent the relations which hold between the parts of the scene. But a change in either the properties or the relations which are part of a perceptible scene will lead to a change in which existential propositions fit it.

Quasi-Descriptivist view makes the relation *R* such that it holds between perceptual experiences and the items they make phenomenally present to their subjects (and holds between no other tuples of items), a highly unnatural relation, in the Lewisian sense of 'unnatural'. Then given that the phenomenal properties of experiences constitutively depend on what those experiences make phenomenally present to their subjects, this may also lead to phenomenal properties themselves coming out as highly unnatural. This would arguably be an unwelcome result.

In fact, however, it is possible to hold the Quasi-Descriptivist view, whilst also holding that there is a single, highly abstract, principle which governs what items an experience makes phenomenally present to its subject in both the successful perceptual and the purely hallucinatory case. In particular, the Quasi-Descriptivist view is compatible with the idea that what an experience makes phenomenally present to its subject is simply determined by the representational content of the experience in question, in both the successful perceptual and the purely hallucinatory case. To see this, note that, given we're already assuming for the sake of the argument that pure hallucinations bear existential content, it doesn't seem much of a leap to assume that successful perceptual experiences bear both *singular and existential* content[101]. Suppose that we take it that the perceptible scene which a perceptual experience makes phenomenally present to its subject is simply the perceptible scene which 'best fits' all of the different sorts of content it carries, in both the successful perceptual and the purely hallucinatory case. Then, in the successful perceptual case, it will be the perceived scene which the successful perceptual experience makes phenomenally present to its subject, since it will both have objects which enter into the experience's singular contents amongst its constituents whilst also, since the

[101] Some philosophers have debated whether the content carried by successful perceptual experiences is singular *or* existential, which suggests that they think these are mutually exclusive options. For orientation in this debate, it is worth reading Tye (2009, ch.4). Tye himself also discusses the pluralist option, which I am relying on here, on which perceptual experiences bear both existential and singular content, though he does not endorse it. For another discussion of pluralism about the representational content of experience see Chalmers (2010, pp.382-4).

experience is a successful perception, fitting the experience's existential content perfectly[102]. Meanwhile, in the purely hallucinatory case, it is merely fit with the experience's existential content which will determine which perceptible scene an experience makes phenomenally present to its subject, for the simple reason that the experience lacks *singular* content altogether[103]. So it looks like we might be able to treat the relation *R* which experiences bear to all and only the items they make phenomenally present to their subjects as a relatively natural relation whilst still getting what we want, by treating it as a relation which experiences bear to the items which best fit the experience's content. This relation doesn't look *especially* unnatural in the Lewisian sense.

Overall, then, the Quasi-Descriptivist view seems to have a lot more going for it than the fact that nothing like it seems to have been discussed in the literature would suggest. In the next section, I will move on to discuss some of the most obvious objections to the view, and show that they are either without merit, or at least depend on highly controversial assumptions which I would not want my argument against Hallucinatory Presence views to rest on.

Before we move to the next Section, however, it is worth noting that there are two different, though not mutually exclusive, ways in which we can look at the argument given in favour of the quasi-descriptivist view in this Section. The first is to see it as simply a piece of normal, empirical, theorizing about the real phenomenon of phenomenal consciousness. On this view, the HMT, the QRM, and extended transparency, are simply things which we know about this phenomenon, and hence, constrain our theorizing about it. The support they give to the Quasi-Descriptivist view and other views on which ordinary objects can be phenomenally present in the purely hallucinatory case, is just the sort of support which being consistent with known facts about a phenomenon

[102] Note that this means that the scene in the world that you are actually perceiving will always be the unique scene present in the successful perceptual case

[103] Technically, this is not quite right if we allow that pure hallucinations carry singular content about themselves or their subjects, but this does not affect the basic point about there being a single principle which says that the (existential and singular) representational content of an experience determines what scene is present that is operative in both the successful perceptual and the purely hallucinatory cases.

lends to a theory. In other words, the support is like that which the class of theories about the atmosphere which are consistent with the sky's appearing blue to normal human perceivers gain from the fact that the sky appears blue to normal human perceivers. An alternative way of understanding what is going on here, is to go meta, and think of the argument given here as part of a project for seeing what phenomenon in the world best fits our concept PHENOMENAL CONSCIOUSNESS[104]. On this view, we can see both the HMT and the QRM, as part of the meaning-fixing theoretical role associated with our phenomenal vocabulary. (Transparency and extended transparency must surely, on the other hand, be empirical theses, given the extent of disagreement about it amongst users of phenomenal vocabulary) We can then see the support which the Quasi-Descriptivist view gains from the fact that it can reconcile these two claims with transparency as being about finding the best candidates in the world for being phenomenal properties.

These two different ways of looking at the argument of this section are not inconsistent, but it is worth keeping them in mind. In particular, if we deny the legitimacy of the second way of looking at things, this will affect the sort of support we think that the HMT and the QRM require. If we think of them as somehow part of the meaning-fixing theoretical role of our phenomenal vocabulary, then they *arguably*[105] (though this is of course controversial) do not need independent support, except support which shows that the relevant theoretical roles can be filled. If, on the other hand, we think of them merely as normal empirical claims about phenomenal consciousness and phenomenal properties which happen to be accepted by a lot of philosophers, then they need further support before we are justified in taking them to be constraints on our theory. I won't return to this issue in the rest of this chapter, but it is worth keeping the two different possible attitudes discussed in this paragraph in mind.

[104] I follow the Fodorian convention of indicating reference to concepts, rather than words, objects, or properties, by capitalization.
[105] Depending on just how Quinean one is about these matters.

Section Three-Objections to the Quasi-Descriptivist View

Here, I go through some of the most obvious objections to the Quasi-Descriptivist view. I won't however, discuss objections that rest upon the denial of the claim that experiences bear existential content, for two reasons. The view that they do is widely shared, and all I'm trying to establish in this Section is that the Quasi-Descriptivist view is not a total non-starter.

On then, to various objections:

1) **Objections from the lack of causal connection between objects and perceptible scenes, in the purely hallucinatory case.**

An obvious first worry to have about the Quasi-Descriptivist view, is that it allows that in the purely hallucinatory case, objects and perceptible scenes can be phenomenally present to the subject of a pure hallucination *H* even though they are not causally connected to *H* in any way whatsoever and nor are any of their constituent objects in the case of perceptible scenes. I have already, in effect, disarmed one form of this objection in the previous section. There we saw that some people might object to any view on which pure hallucinations make objects and perceptible scenes present to their subjects on the following grounds. There is no causal connection[106] between pure hallucinations and any objects or perceptible scenes. So there must be no good answer to the question which such object or scene a particular pure hallucination makes phenomenally present to its subject. We saw, of course, that the Quasi-Descriptivist view easily has the resource to deal with this objection, since we can treat the existential content of the experience as determining which perceptible scene is present, and hence which objects. However, there are other versions of the objection which I did not directly respond to in the

[106] Or at least no causal connection 'of the right kind'

previous Section. In particular, we might worry that there's something wrong at some kind of conceptual level with the very idea that an experience could make phenomenally present to its subject particular objects, without those objects bearing the right sort of causal relation to that experience.

However, there doesn't seem to be any reason to think that there's any conceptual connection (in any sense of that vexed idea) between an experience's making items phenomenally present to its subject, and those items standing in some sort of perceptual or causal relation to the subject of that experience. Recall that, when I introduced the notion of 'phenomenal presence' in Section One, I did so largely as a placeholder for talk about the items associated with a successful perception *E* which are the best candidates to be *E*'s qualia. The only independent thing I said to specify what 'phenomenal presence' is supposed to pick out was that it was connected to a certain sort of metaphorical talk about our perceptual experiences which philosophers have tended to find appropriate. This was the metaphor of 'presence', the idea that when we perceive (or hallucinate) items those items seem to be somehow 'within' the perceptual experience, in a way in which the items which occurrent thoughts are about do not seem to be 'within' those thoughts. None of this suggests that there's any sort of 'conceptual' connection between an experience's making an item phenomenally present to its subject, and that item standing in any sort of causal relation to an experience. There's clearly no conceptual connection between being an item which is one of the best candidates for being the qualia associated with an experience, and standing in any sort of causal relation to that experience. After all, the traditional view of qualia is that they themselves are properties of experiences, and the properties of a perceptual experience *E* are not involved in causing *E*. And a metaphor which simply functions to point at a phenomenon that we know when we see it, even though we cannot articulate any non-metaphorical description of it, can hardly be the source of much in the way of conceptual constraints. Perhaps the relevant constraints come from how the phenomena characterized metaphorically must be for the metaphor to be appropriate. But even if there are such

constraints here, it's hard to think why they'd imply that only causes of experiences are made present by those experiences to their subjects.

2) The Quasi-Descriptivist View Fails to Secure Stronger Formulations of the HMT, and it is Really From These That the Weak Version I Have Taken as Canonical Derives its Support

A second worry for the Quasi-Descriptivist view is that although it can secure the truth of the HMT, there will still be some possible pure hallucinations which ought to match successful perceptions and yet do not do so.

In particular recall that in ch.1 the HMT was introduced as the weakest and least controversial formulation of the intuition that whether an experience *E* is a pure hallucination or a successful perception does not, by itself make any difference to its phenomenal properties. But I also said there that it's not clear that the HMT is the best possible formulation of this intuition. We saw in ch.1 that there are other possible ways to formulate the HMT, and that it may well be the case that one of these other formulations better captures our intuitive commitment that, in some sense, whether or not an experience is a pure hallucination or a successful perception should not, in itself, make a difference to its phenomenal properties. In ch.1, I suggested three alternative ways of formulating the HMT, of increasing orders of strength. Firstly in terms of nomological, rather than metaphysical possibility. Then in terms of it being the case that, for any successful perception, *E*, counterfactually, were one to turn *E* into a pure hallucination be destroying the perceived objects in the world, but keeping its neural base intact, *E*'s phenomenal properties would not change. And finally in terms of the phenomenal internalist claim that no two experiences can differ phenomenally if they are narrow duplicates and occur within subjects who are narrow duplicates over their respective life-times.

Now, if we formulate the HMT in terms of nomological, rather than metaphysical possibility, I do not think that there is any problem here for the HMT. Plausibly, there will be, for each successful

perception *E*, a nomologically possible world, *w*, which contains both a pure hallucination *H* which shares *E*'s existential content, a scene *S* which is the only scene which exactly fits that existential content at that world. If so, the HMT will come out as true on the Quasi-Descriptivist view, even if we formulate it in terms of nomological rather than metaphysical possibility. Difficulties do however, follow, for the Quasi-Descriptivist view, if we take it that our intuition is best captured by the counter-factual, or phenomenalist internalist claims.

Suppose that our intuitive commitment to the idea that it ought not to make a difference to an experience's phenomenal properties whether it is a pure hallucination or a successful perception, really is best formulated in terms of one or other of these claims. And suppose the Quasi-Descriptivist view although it does guarantee that the HMT itself is true, is also a view on which one of these stronger extensions of the HMT come out as false. If so, it might be in one or other of two kinds of trouble.

Firstly, if we are committed to one (or both) of these two stronger principles, then the fact that the Quasi-Descriptivist view is inconsistent with that principle would itself count against it.

Secondly, suppose that we only believe the HMT because we are committed to one (or both) of the two stronger claims, and have no other strong evidence for the HMT. If the Quasi-Descriptivist view turned out to be a view on which that stronger claim was false, its being consistent with the HMT itself would no longer count in its favour. Instead, in order for our background commitment to the idea that whether an experience is a successful perception or a hallucination by itself tells us nothing about its phenomenal properties, to give us evidence for the Quasi-Descriptivist view, the latter would have to be consistent with whichever of the stronger extensions of the HMT captures our real commitments here. And it's not clear the Quasi-Descriptivist view entails the counter-factual or the phenomenal internalist claims. The Quasi-Descriptivist view as we've seen, plausibly guarantees the following: for any successful perceptual experience, *E*, which makes present to its subject some perceptible scene, *S*, there is, somewhere

in modal space, a possible world containing a pure hallucination, *H*, which makes present to its subject a scene qualitatively identical to *S*. But this falls well short of showing that *all* metaphysically possible pure hallucinations which share *E*'s narrow properties make such a scene phenomenally present to their subjects. It also falls short of showing that, in the nearest world where someone keeps *E*'s neural base intact, but destroys *S*, making *E* into a pure hallucination, the resulting pure hallucination would make a scene qualitatively identical to *S* phenomenally present to its subject.

In fact, however, it turns out to be hard to get an especially convincing objection to the Quasi-Descriptivist view out of the thought that we might be committed to the HMT only because we are committed to the counterfactual, or the phenomenal internalist, claim. The Quasi-Descriptivist view is clearly going to be inconsistent with phenomenal internalism. On the Quasi-Descriptivist view, the phenomenal properties of an experience are determined by the perceptible scene it makes present to its subject. And the latter is determined partly by the existential content of the experience, but also partly by what scenes exist. But clearly, two possible experiences which are narrow duplicates can differ with respect to the range of perceptible scenes they have as world mates. So the phenomenal properties of narrow duplicates will be able to vary on the Quasi-Descriptivist view, contrary to phenomenal internalism. But it's not clear how troubling this is for the advocate of the Quasi-Descriptivist view, because it turns out to be *highly* controversial that the phenomenal internalist claim is a good way of formulating the basic intuition behind the HMT. Meanwhile, if we choose the counterfactual formulation of the basic idea behind the HMT, it turns out to be far from obvious that this counterfactual claim will come out false on the Quasi-Descriptivist view, even if we also have the QRM and extended transparency[107].

[107] The Quasi-Descriptivist view is only a Hallucinatory Presence view, and hence only of interest if it allows us to preserve the latter two claims.

Firstly, on phenomenal internalism. Phenomenal internalism is simply too controversial as a gloss on the intuition behind the HMT for it to be the case that the Quasi-Descriptivist view does not really help us reconcile that intuition with extended transparency and the QRM if it does not provide us with a way to reconcile the HMT and the QRM with phenomenal internalism. Phenomenal internalism is incompatible with the phenomenal properties of experiences being able to vary with changes in their wide content, assuming their narrow content does not also vary. Yet some of the most prominent advocates of representationalism, who also endorse the HMT, such as Michael Tye, think that changes in the wide existential content of experiences entail changes in their phenomenal properties. Since some prominent advocates of the HMT deny phenomenal internalism it would be dubious to rely on the premise that the intuition behind the HMT is best captured by phenomenal internalism.

Things are much more complicated with the claim that the intuition behind the HMT is best expressed via the claim that for any successful perception *E*, were one to destroy the scene *E* is a perception of whilst keeping *E*'s neural base intact, its phenomenal properties would not change. Call this claim CHMT (for counter-factual HMT). I think CHMT may well be false, given the Quasi-Descriptivist view (plus the QRM and extended transparency). But it may also come out true. Surprisingly, it depends on our background assumptions about cosmology.

CHMT specifies a set of counterfactuals which must be true. Counterfactuals are true (at least on the Lewis-Stalnaker semantics[108]), iff in the closest world(s) where the antecedent of the counter-factual is true, the consequent is true. This means that the following must hold if the CHMT is true. It must be the case that for each (actual) successful perceptual experience, in the nearest world(s) where the perceptible scene associated with that experience is destroyed, but the neural base of the experience is itself preserved, the phenomenal properties of that experience don't change. Now, it may seem very unlikely that such counterfactuals come out true very often, let

[108] See Lewis (1973)

alone that they come out true in all cases. After all, how likely is it that, at the world(s) most like the actual world where some actual successful perception *E* is turned into a pure hallucination by the destruction of the perceptible scene, *S*, which it makes present to its subject, there will simply be a carbon copy of *S*? Suppose that *E* is a perception of a mountain, *M*. How likely is it that there happens to be a mountain which has exactly the properties that *M* appears to have, somewhere else in the universe?

In fact, however, it's not at all obvious that it will always or usually be the case that at the nearest worlds where *E* is turned into a pure hallucination, there will be no perceptible scene qualitatively identical to that which *E* actually makes present to its subject. For, it turns out that there are live hypotheses taken seriously in contemporary cosmology which some leading physicists have argued imply that in the actual world every nomologically possible event that can happen does happen and indeed happens infinitely many times. In particular, Tegmark (2014, ch.s5-6) discusses the fact that many cosmologists today believe that the universe is infinitely large. Tegmark claims that it follows from this, plus certain features of quantum mechanics, that 'everything that can happen does happen', or in other words that all events consistent with the laws of physics occur somewhere in our universe. Indeed, Tegmark claims that it follows from reasonably mainstream cosmological theory that not only does every nomologically possible event occur, but it occurs infinitely many times.

If Tegmark turns out to be right about all this, it presumably follows that every nomologically possible perceptible scene exists somewhere in our universe. If this is so, then presumably in the *nearest* possible worlds where particular successful perceptual experiences are changed into pure hallucinations by the destruction of their associated perceptible scenes, there will in fact be a qualitatively identical scene somewhere else in those worlds. For, since they were successful perceptions, the perceptible scenes which they made present to their subjects are certainly nomologically possible. And since, for each successful perception the nearest world where it

becomes a pure hallucination will presumably have a cosmology similar to ours, then there will be infinitely many scenes within that world which are qualitatively identical to the scene which the successful perception makes present to its subject. Assuming that the pure hallucination shares the existential content of the original successful perception, those qualitatively identical scenes will exactly fit the existential content of the pure hallucinations. Now, suppose that we have a version of the Quasi-Descriptivist view on which, when there are multiple perceptible scenes which exactly fit the existential content of a pure hallucination, it is vague which of those scenes it makes present to its subject, but determinately true that it makes a scene of the (maximally determinate) qualitative type that all those scenes fall under present to its subject. Then the relevant counterfactual will come out true. For it will be (determinately) true at the nearest possible world where a successful perception, *E* is changed into a pure hallucination, *H*, by destruction of the objects of perception and preservation of its neural base, that *H* makes present to its subject a qualitatively identical scene to the scene which *E* makes phenomenally present to its subject. This will be true despite there being no particular scene, *S*, such that it's determinately true that *H* makes *S* phenomenally present to its subject[109]. Admittedly, this does narrow the space of possible versions of the Quasi-Descriptivist view, in that someone who is opposed to the idea that it can be vague which particular perceptible scene a pure hallucination makes present to its subject will now have no direct reply to the objection from the CHMT, but I would not want to rest my case against Hallucinatory Presence views on this assumption. Equally, I would not want to rest it on the assumptions that a particular live hypothesis in cosmology is wrong, or that Tegmark is wrong about what it implies.

[109] Though this does depend on assuming that when one experience *E* makes present to its subject a particular perceptible scene *S*, and there is another experience *E*' of which it is determinately true that it makes present to its subject a scene qualitatively identical to *S*, but there is no scene *S*' such that it is determinately true that that *E*' makes it phenomenally present to its subject, *E* and *E*' match phenomenally. But this looks like the sort of fine point which we ought to let our best theory of the metaphysics of phenomenal consciousness decide, rather than taking an independent stance on it.

The above, I think, is enough to disarm any argument against the Quasi-Descriptivist view from the CHMT, especially since it's not especially clear that the CHMT is the best way of formulating the intuition behind the HMT in the first place.

Section Four-Two Genuine Problems for the Quasi-Descriptivist View

In this section, I present an objection to the Quasi-Descriptivist view which I think raises more genuine problems for it. The objection is not, I think obviously fatal to the view. Indeed I would judge it to be weaker than the argument of ch.2.3 is as an objection to *property-based* Hallucinatory Presence views. At the same time, the objection does not strike me as one to which the proponent of the Quasi-Descriptivist view has an obviously adequate response. Hence I think the objection is at least strong enough to make it worth looking at alternatives to the Quasi-Descriptivist view, despite the fact that it does not clearly show it to have unacceptably high costs. After discussing the objection, I very briefly discuss the prospects for views other than the Quasi-Descriptivist view on which pure hallucinations can make perceptible scenes phenomenally present to their subjects.

The objection is that it's unclear how the Quasi-Descriptivist view can be extended to cover illusions and impure hallucinations. Recall that, as well as (fully) successful perceptions and pure hallucinations, there are cases of partial perceptual success; cases of illusion where objects are perceived but perceived to have properties they don't in fact have, and cases of impure hallucinations where objects are hallucinated at the same time as other objects are perceived. If the Quasi-Descriptivist view is to be a full account of the metaphysics of perceptual consciousness, it needs to be able to tell us what is going on in such cases of impure hallucination

and illusions. From now on I will refer to these collectively as semi-successful perceptions. But there are difficulties here for the Quasi-descriptivist view.

In particular, the most obvious way to extend the Quasi-Descriptivist account in these cases, is simply to say that in the semi-successful case, as in the successful perceptual or the purely hallucinatory case, experiences make present to their subjects the perceptible scenes which best fit the representational content which they bear. The phenomenal properties of illusions and impure hallucinations would then be fixed by which items they make phenomenally present to their subjects, in just the same manner as they are in the (fully) successful perceptual case, or the purely hallucinatory case.

Unfortunately, however, an awkward question arises about this proposal in the semi-successful case, which is not troublesome in the other two cases.

The question is 'how do we weigh fit with the existential content of an experience against its singular content in determining which perceptible scene best fits the existential content of an experience'? This is not a bothersome question in the (fully) successful perceptual case. It is obvious that the perceptible scene which best fits the singular content of a successful perception will be the scene *S* consisting of the perceived objects instantiating the properties they are perceived to have and related the way they will be perceived to be. Further *S* will also fit the existential content of the experience at least as well as any other. After all, *S* just consists in those objects related to each other in exactly the way that the singular proposition associated with the experience in question says they are, and instantiating the properties that it says they instantiate. Furthermore, *S* will also perfectly fit the existential content of the experience, since the experience's existential content just is the singular content with all reference to particular objects quantified out. Since in the successful perceptual case, there will always be a particular perceptible scene which is the unique best fitter of the singular content, whilst fitting the existential content no worse than any other scene, we don't need to have a way of weighing fit

with the existential content against fit with the singular content. As for the purely hallucinatory case, either there are no singular propositions associated with purely hallucinatory experiences, or if there are, then they are propositions of an odd, gappy type, in which the only actually singular element involves the subject of the experience (or the region of space in which the subject or the subject's head is located perhaps.) On the first understanding, there is no issue about weighing the singular content against the existential in the purely hallucinatory case, since there is no singular content. On the second view, if we can make sense of fit with 'gappy' singular content, then there is an issue about how to weigh this against the existential content of the experience. But since the only singular element in the gappy content involves the subject of the experience, this issue does not look too serious on its own. With the singular element in such cases being so minimal, we would arguably be keeping to the spirit of the Quasi-Descriptivist view if we simply used it to break ties between perceptible scenes which otherwise fit the existential content of the experience exactly equally well[110].

Unfortunately, however, whilst in the semi-successful case the experience will always have some singular content into which objects in the external world enter, there is no obvious guarantee, unlike in the (fully) successful case, that the perceptible scene which best fits the singular content of the experience will also fit the existential content of the experience better than any other scene. Intuitively, this is because how well something fits the existential content of an experience depends only on how well it fits a particular qualitative description, but how well it fits the singular content presumably depends at least partly on whether it contains the objects which enter into the singular content as parts. From this it presumably follows that it's possible for which perceptible scene best fits the existential content of an experience to come apart from which best fits its singular content.

[110] I already suggested this manoeuvre in Section Two.

Suppose, for example, that Tegmark is right and that the universe really does contain every nomologically possible perceptible scene. (The assumption is harmless since such a universe is surely possible and the Quasi-Descriptivist view is presumably necessarily true if true at all.) And suppose that a subject, *S*, undergoes an experience *E*, which is a genuine perception of a red square *r*, but which is very slightly illusory with respect to *r*'s exact shade. Since we're in a universe which contains every nomologically possible perceptible scene, there will be a perceptible scene which has as a part a square which is exactly the shade of red which *r* appears to be to *S*, and which fits the existential content of the experience better than any perceptible scene containing *r* as a part. But it's not obvious that such a scene would also fit the singular content of *E* better than any perceptible scene containing *r* as a part. Since *r*, the particularly perceived square itself, enters into the singular content of *E*, presumably fit with that singular content is partly determined by whether or not *r* is part of a particular perceptible scene. Since whether or not a perceptible scene contains *r* as a part matters for fit with *E*'s singular content, but not its existential content, which perceptible scene best fits one content can come apart from which perceptible scene best fits the other.

Now, if fit with the singular content of an experience really can come apart from fit with the existential content, this creates problems for the Quasi-Descriptivist view. The Quasi-Descriptivist view was originally introduce recall, to disprove the claim that there is non-arbitrary way of associating particular perceptible scenes with pure hallucinations. It did so by suggesting a reasonably natural relation which our use of phenomenal vocabulary might have glommed onto which did take pure hallucinations to specific perceptible scenes (and objects). However, if which perceptible scene best fits the singular content of a semi-successful perception can come apart from which best fits its existential context, then the threat of indeterminacy will return in a slightly new guise in the semi-successful case. For, if there are cases where the perceptible scene, S_1 which best fits the *existential* content of an experience *E*, is not identical to the perceptible scene, S_2 which best fits *E*'s *singular* content, then we need a way to weigh the importance of fit

with the existential content of an experience against fit with its singular content in determining which of these scenes *E* makes present to its subject[111]. And it's not clear how there could be a non-arbitrary way of determining how much weight to give to the singular content of an experience, as against its existential content, in determining which perceptible scene best fits it. Nor is it clear exactly what it would mean to set a particular weighting here. There's certainly no obvious way of making sense of the idea of, say, treating fit with the existential content of an experience as being twice as important as fit with the singular content.

Incidentally, similar problems emerge when we try to work out the idea of best fit with the singular content alone, in the semi-successful case. We can treat fit with singular content in a variety of ways. For example, are all perceptible scenes which have an object which enters into a singular content amongst their parts 'better fits' for that singular content than all other perceptible scenes? Or can fit with the qualitative aspect of that singular content sometimes outweigh this? If we answer yes to the latter question, then how do we trade off fit with the qualitative aspect of the singular content of an experience as against the fact that a perceptible scene contains objects which enter into that singular content? Once again these questions are hard to answer.

The problem here is that, unless we can come up with a good way to answer these questions about how to weigh the singular content of an experience against its existential content, and what it is to fit the singular content of an experience to a degree (rather than fully or not at all), then there will remain a worry that any particular answer we give to these questions will be arbitrary. If answers to these questions are arbitrary however, then this raises a worry that our use of phenomenal vocabulary won't really pin down a very precise meaning for 'phenomenally present to', at least in the semi-successful case.

[111] Unless we are prepared to say that it is determinately true that it makes one or the other scene present, but vague which.

Note that the problem here is not that there are no relations which take illusions and impure hallucinations to perceptible scenes. Provided we are assuming a plenitudinous conception of properties and relations, there will certainly be such relations, and amongst them will be relations which also get the successful perceptual and purely hallucinatory cases right. Rather the problem is that there will be many, many such relations. And all of these relations will disagree on which perceptible scenes go with which semi-successful experiences, in those cases where there is not an obvious 'winner' for best fit[112]. This makes it hard to see how we could be glomming on to any one of these relations. It also makes it harder, I think, to see the Quasi-Descriptivist view as one on which 'makes phenomenally present to' comes out as reasonably natural in the Lewisian sense. On some metasemantic views, this will lessen the chance that it is the relation of best fit with an experience's content that our phenomenal presence talk is glomming on to.

It's unclear to me just how troubling the problem here is. One possible response is to hold that things go vague here, and that, in cases where best fit with the existential content comes apart from best fit with the singular, or where it's not clear what best fit with the singular amounts to, it's vague which perceptible scene is phenomenally present to the subject of the experience. After all, we saw in Section Two that unless we hold that pure hallucinations in worlds with more than one scene which perfectly fit their existential content make *no* scenes phenomenally present to their subjects there will be some vagueness even in the *purely* hallucinatory case. So if someone thinks that what is phenomenally present is always determinate, or that the phenomenal properties of an experience are always determinate[113], they will already be off the boat with the Quasi-Descriptivist view anyway. Still, the vagueness here feels more troubling, in that it doesn't seem a product of a particular, fully-worked out idea of best fit with a description producing ties in some case, but rather a product of us not really having a proper grip on what 'best fit' with the content of an experience is, in the semi-successful case. This is why the

[112] Here by an obvious winner, I mean a scene which fits both the singular and the existential content best on *any* reasonable way of working out the idea of 'fit' with a description.)
[113] What is present fixes the phenomenal properties, recall.

vagueness here threatens the idea that there's an obviously natural joint for the meanings of our phenomenal vocabulary to get sucked towards.

Of course, I haven't *proven* that there is no principled way to determine which scenes are present in the semi-successful case, available to the proponent of the Quasi-Descriptivist theory. Perhaps there is some way of explicating the idea of best fit with the content of an experience in the semi-successful case which is in fact obviously the one to which the meanings of our terms will be drawn towards. But it's certainly not obvious to me how this would work.[114]

So the Quasi-Descriptivist view faces at least one reasonably serious looking problem with semi-successful perceptions. What about other views on which perceptible scenes containing ordinary physical objects as parts can be phenomenally present to us in the purely hallucinatory case? Well, recall that all such views do face the problem of saying *which* perceptible scenes pure hallucinations make phenomenally present to their subjects. It was in response to this problem that we suggested that perhaps pure hallucinations make phenomenally present to their subjects the scenes that best fit their existential representational content. Now we have just seen that it is hard to see how to apply this idea to semi-successful perceptions. But if we drop the idea, then we are left once again with the original problem about how it gets determined which perceptible scenes pure hallucinations make phenomenally present to their subjects. Furthermore, if we are deprived of the resource of using the existential content of pure hallucinations to explain this, it's not immediately clear where an alternative explanation can be found. I therefore tentatively

[114] Incidentally, there is another potentially serious problem for the Quasi-Descriptivist view. Some philosophers seem to think it is *a priori* or at least obviously true, that the only properties which experiences make present to their subjects are those which enter into their representational content. At least, I take this to be a view held by advocates of the 'phenomenal intentionality' thesis, such as Horgan and Tienson (2002), who believe that the facts about (to put things in my terms) what items a perceptual experience makes phenomenally present to its subject, fix the facts about what properties it represents there to be instantiated in the world. I also take it to be held by those representationalists such as Byrne (2001), who argue for representationalism on the grounds that any change in the properties made phenomenally present to us by an experience also marks a change in how the experience represents the world to be. I think that this principle is likely inconsistent with the quasi-descriptive view, for reasons I don't have space to go into here. However, the view I tentatively put forward as my own in the next chapter also rejects the putatively obvious/*a priori* principle cited here, so it would be dialectically inappropriate.

conclude that the consensus view that pure hallucinations can't make particular physical objects, or perceptible scenes which have particular physical objects as parts, phenomenally present to their subjects is likely correct. If so, then the argument of ch.2.3, does indeed give us good reasons to reject all Hallucinatory Presence views, not just the property-based ones. Likewise, if the only viable Hallucinatory Presence views are property-based ones, then the fact that property-based views are inconsistent with nominalism, is a problem for anyone who wants to endorse a Hallucinatory Presence view.

I tentatively conclude, then, based on the argument of this and the previous chapter, that we should reject Hallucinatory Presence views. In the next chapter, I discuss whether we should instead adopt a Phenomenal Difference view or an Anti-Presence view, arguing for the former. Note however, that I would not be too upset if someone else weighed the evidence differently, and decided to prefer the quasi-descriptive view to the view I eventually endorse in the next and last chapter. This looks especially likely to be reasonable if they either were confident that the intuition behind HMT is best formulated in terms of metaphysical possibility, rather than anything stronger, or are confident that ours is a universe where every nomologically possible event occurs. The view I defend in ch.4 is certainly not free of difficulties. It's therefore, to a certain extent, a judgment call, as to whether those difficulties are more severe than those the Quasi-Descriptivist view faces.

Chapter Four: A New Hybrid Naive Realist-(Almost) Representationalist View of Phenomenal Consciousness

Introduction-

In the previous two chapters, we've seen that there are good reasons for rejecting Hallucinatory Presence views. In particular, we saw that, an extended version of the transparency thesis is highly plausible, according to which our experiences make perceptible scenes, which have ordinary objects as parts, phenomenally present to their subjects. We've seen that, if we combine this with the most plausible formulation of the idea that two experiences differ in their phenomenal properties if there is a difference in the qualia-role playing items associated with them, the QRM, we get the following result. Pure hallucinations can only match successful perception if they also make perceptible scenes which have ordinary objects as parts phenomenally present to their subjects. In the last chapter, we saw that it's hard to come up with a good account of how pure hallucinations could do this. Hence, there doesn't seem to be a way to hold on to all three of transparency (in its extended) form, the QRM and the HMT. In this chapter, I will discuss whether we should adopt a Phenomenal Difference view, and drop the HMT, whilst keeping the QRM and (extended) transparency, or whether we should instead adopt

an Anti-Presence view, and drop the QRM, whilst keeping the HMT and (extended) transparency[115].

This chapter has four sections. In Section One, I argue that we should prefer a Phenomenal Difference view to an Anti-Presence view. In Section Two I then develop a particular Phenomenal Difference view, which combines elements of naïve realism and representationalism, whilst hopefully avoiding the worst problems of either. In Section Three, I show that my view is not vulnerable to standard arguments against either naive realism or representationalism. Finally, however, in the last section of this chapter, I show that my own preferred Phenomenal Difference view encounters difficulties *if* illusion turns out to be *extremely* pervasive. So the precise Phenomenal Difference view I defend in this chapter should be seen as the view, I endorse conditional on two other claims. Firstly that the problems with the Quasi-Descriptivist view (and other views on which perceptible scenes can be present in the purely hallucinatory case) discussed in ch.3 can't be solved. And secondly, that illusions are not extremely pervasive, in the manner discussed in Section Four. I discuss what other view we should endorse if illusion is extremely pervasive in Section Four. I do not come to any firm conclusions. But one of the suggested views is a merely slightly modified version of the particular Phenomenal Difference view defended in Sections Two and Three.

Section One-Anti-Presence Views versus Phenomenal Difference Views

Here, I want to give some reasons for preferring Phenomenal Difference views, which keep transparency (and extended transparency) and the QRM, but reject the HMT, over Anti-Presence views, which reject the QRM, but keep the HMT, along with transparency. I won't examine the

[115] As I said in ch.1, I won't discuss the option of dropping transparency (or extended transparency), for space reasons.

prospects of views which try to weaken the QRM, whilst keeping (extended) transparency, and the HMT. Such a view, would be a view on which some similarities in what two experiences, E_1 and E_2 make phenomenally present to each other is a necessary condition on their matching phenomenally, but the similarity required is weaker than there being no qualitative difference in the items they make phenomenally present to their respective subjects. As I said in ch.2, it's hard to find a *principled* weakening of the QRM which both preserves the basic idea that the qualia role players associated with an experience fix its phenomenal properties, whilst also accounting for cases where differences in the structured items experiences make present to their subjects make for a phenomenal difference between those experiences. This is why I take Anti-Presence views, i.e. views which reject not just the QRM, but any view on which two phenomenally conscious experiences matching is a matter of there being sufficient similarities as the only plausible rivals to Phenomenal Difference views, other than Hallucinatory Presence views.

Why prefer a Phenomenal Difference view to an Anti-Presence view, then?[116] Well, my basic case is that the main relevant difference between such views is that the former allow us to preserve the QRM, but not the HMT, whilst with the latter the reverse is the case. I think that the case for the HMT is weaker than the case for the QRM. We saw in ch.1, both that the HMT is very widely accepted, with its denial felt to be 'counter-intuitive', but also that the QRM (or at least something very like it) is built into the standard account of phenomenal properties. So we have *prima facie* reasons to endorse both. But I will argue that our reasons to endorse the HMT are in fact weaker than our reasons to endorse the QRM.

[116] Some might worry that, since the QRM (or something like it) and the HMT are both in some sense 'part of our concept of phenomenal properties', the debate here goes verbal. Three points in response to this worry. Firstly, insofar as I give reasons below for preferring one to the other, this makes the debate look less verbal. Secondly, given that phenomenal properties arguably play an important role in ethical theorizing, it may matter which of the views is correct, even if the dispute is verbal, especially if we hold a metaethical view on which things, including phenomenal conscious experiences have value (or disvalue) *because* we value them. Finally even if the dispute between Phenomenal Difference views and Anti-Presence views is both verbal, and unimportant because verbal, the rest of this chapter may be of interest in its attempt to work out what the best Phenomenal Difference view looks like. The latter may not be a verbal matter, even if the former is.

I offer three arguments for this.

Firstly, as we saw in ch.1, it's not clear how best to formulate the basic intuition behind the HMT. This, I think, should cause us to doubt that we really have much grip on the supposedly obvious truth which it is supposed to be a precise statement of. Secondly there is reason to be suspicious that we are only committed to the HMT because we are committed to phenomenal internalism. This is problematic for several reasons. Phenomenal internalism is itself somewhat controversial, and explicitly rejected by some proponents of the HMT (Tye (2009, ch.8)). It requires somewhat controversial assumptions to reconcile transparency with phenomenal internalism. And most importantly it's very hard to reconcile phenomenal internalism with extended transparency. Thirdly, I will argue that we should distrust our commitment to the HMT because I suspect our commitment to it is a result of a background commitment to the phenomenal presence of the same items in the purely hallucinatory and the successful perceptual case. This suggests that we are only committed to it because we are committed to a larger picture of what's going on in the purely hallucinatory case. Furthermore, we saw in ch.s2 -3 that this picture is false; successful perceptions make perceptible scenes present to their subjects and pure hallucinations don't. If this is right, then the HMT should surely be assigned less weight than the QRM, given that our commitment to it is rooted in a false picture of what pure hallucinations are like.

I don't think any of these three arguments are individually decisive. But I think that together they give us decent reason to think that the case for the HMT, given our overall evidence, is weaker than the case for the QRM. If so they make a decent case for favouring Phenomenal Difference views over Anti-Presence views. Note however, that even if I'm wrong about that, the development of my own preferred Phenomenal Difference view in the remainder of this chapter may not be totally without interest. It may be that, though the arguments I have given do not establish a reason to prefer Phenomenal Difference views to Anti-Presence views, still we do not yet have good reasons to prefer the latter to the former. If so, then at least my development of

my own preferred view will be a development of one of the serious contenders for being the right type of view.

Let's proceed then, to the three arguments:

1) Recall that in ch.1, I noted that the HMT is not obviously the best way of making precise the basic idea that whether or not an experience is a successful perception need not make any difference to its phenomenal properties. Setting aside the issues of how to formulate the HMT in a way which covers impure hallucinations and illusions, as well as pure hallucinations, we saw in ch.1 that the claim that for any successful perception, there 'can' be a pure hallucination which shares all and only its phenomenal properties, can be interpreted in very different ways. We saw, recall, that this claim might interpreted as a claim about what metaphysically possible pure hallucinations there are (i.e. those which match successful perceptions), or as a claim about what nomologically possible hallucinations there are, or as a counter-factual claim, or as ultimately amounting to a commitment to phenomenal internalism. Now, this multiplicity of ways of cashing out the HMT ought already to make the defender of an Anti-Presence view a *little* nervous. If we cannot even agree on exactly what the HMT is, then why think that we can be so sure that it is true? Still, this is at most a minor worry.
2) Much more serious is the fact that, if it turns out our commitment to the HMT is generated by a background commitment to phenomenal internalism trouble follows. Phenomenal internalism of this sort will certainly entail the HMT, since there are definitely pure hallucinations which are narrow duplicates of successful perceptions. Furthermore, I already mentioned there, that denial of this kind of phenomenal internalism is often claimed in the literature to be highly counter-intuitive[117]. So there is

[117] Those who claim phenomenal internalism of the kind I'm discussing here to be obviously true or at least highly intuitive include Hawthorne (2004), (Chalmers , 2010, p.389), Kriegel (2007, p.321), all cited in this context by Pautz (2014). Those who themselves deny phenomenal internalism, but still claim that its denial

evidence that many philosophers at least, are committed to it, which at least could be what explains the fact that many philosophers endorse the HMT.

This is troublesome for anyone who wants to use the HMT to support Anti-Presence views. Despite often being consider to be somehow obvious or intuitive, Phenomenal internalism, as I mentioned above, is quite controversial. And we'll see below that it faces some reasonably serious problems. If our commitment to the HMT depends on a background commitment to a controversial and problematic view, this would presumably weaken our reasons for endorsing the HMT itself. Hence it would weaken the degree of support which Anti-Presence views gain from consistency with the HMT.

Firstly, most well developed Anti-Presence views are representationalist views. And Anti-Presence views are also partly defined as views that are compatible with transparency. It is difficult, however, to combine a commitment to representationalism and transparency with phenomenal internalism, without ending up committed to the somewhat controversial thesis that there is narrow as well as wide content.

Transparency, recall, says that when we try and introspect our phenomenally conscious perceptual experiences, in the successful perceptual case we find ordinary physical objects and/or their properties (and do not find any mental items). By far the easiest explanation of this is that the items we find when we introspect our experiences in the successful perceptual case are simply the items our experiences represent. Proponents of Anti-Presence views by definition reject the claim that across all cases two experiences match phenomenally only if there's no qualitative difference in what they make present to their subjects. But it seems likely that they will want to allow that *in the successful perceptual case*, qualitative changes in the items an experience makes phenomenally present to its subject entail changes in the phenomenal properties of the corresponding experiences, and vice versa. Without some claim like this it's unclear why we should take

is widely thought to be counter-intuitive include Devries (1996), Dretske (1995, p.125), Ellis (2010), Horwich (1996) and Tye (2009, ch.8.6).

Anti-Presence theorists to be talking about phenomenal properties at all, given the importance of qualia talk in fixing the meaning of our phenomenal vocabulary. If so, it seems like proponents of representationalist Anti-Presence views will want to acknowledge that differences in what properties enter into the representational content of a *successful perception* always make a difference to its phenomenal properties. But now it looks like we can preserve phenomenal internalism only if we believe that there is narrow content. Otherwise, it looks like in Twin Earth style cases, we will have differences in what properties two successful perceptions represent without differences in their relevant phenomenal properties, given phenomenal internalism. This is in conflict with the claim that differences in at least existential content *between successful perceptions* entail differences in their phenomenal properties. So it looks like the best developed type of Anti-Presence views which are compatible with transparency, representationalist views, will, on plausible further assumptions, entail that there is narrow content. But it's somewhat controversial whether there is any such thing.

More seriously, there appears to be a conflict between phenomenal internalism, Anti-Presence views, and *extended* transparency. There, recall, I argued that amongst the items phenomenally present to us in the successful perceptual case are perceptible scenes. Now, which perceptible scenes are available at a particular possible world is obviously an eminently external and contingent matter, not especially connected to what is going on in the brain of any one subject[118]. So, if our experiences make perceptible scenes phenomenally present to us, then there can definitely be pairs of *successful perceptions* with all the same narrow properties which differ in the perceptible scenes they make phenomenally present to their subjects. Yet presumably, as I said above, even proponents of Anti-Presence views will want to think that the phenomenal properties of *successful perceptions* change with dramatic qualitative changes in what is phenomenally

[118] In comparison, non-extended transparency only makes a claim about properties, which may well be necessary existents, and so is not obviously subject to this argument.

present to us. This strongly suggests we should not appeal to the HMT to motivate Anti-Presence views, *if* our commitment to the HMT itself is really a product of phenomenal internalism.

3) My last argument against Anti-Presence views also involves the worry that we are committed to the HMT only because we are committed to something more problematic that entails it. This time, the worry is that we are really inclined to endorse the HMT, because we are inclined to endorse the following idea. Roughly speaking, what cases of (visual) hallucination *just are* is cases where things are the same within some kind of phenomenal visual field as they would be if the subject had a corresponding successful perceptual experience, but no object in the world plays a causal role in bringing about how things are in the phenomenal visual field.[119] I suspect that many philosophers also tacitly hold a similar picture of illusions. Note that this is a different idea from the phenomenal internalism discussed in the last section, although there might be a connection between them, if only because it is easy, I think psychologically, to pass from the sort of picture just mentioned to phenomenal internalism. In itself though, the idea of hallucinations and illusions as cases where what's going on in some kind of visual field remains the same, but the world differs, does not actually *entail* that whether or not a subject is hooked up to the entities which inhabit the visual field depends only on what goes on in their heads[120].

   Now, if this is the reason we accept the HMT, at least if it is the principal or only reason, this ought to make us doubt it. If the argument of ch.s2.-3, is correct, then successful perceptions make perceptible scenes phenomenally present to their subjects, but no

[119] Several parts of this claim could do with a lot of unpacking if I was seriously putting it forward as my view, but instead I'm just trying to capture a vague, inchoate and tacit idea which I suspect might be lurking in the back of people's minds, when they find the HMT to be obviously true. For example, the claim would have to be reformulated to take account of the possibility of deviant causation of pure hallucinations by perceptible objects.

[120] It's compatible, for example, with the Tye-Dretske style view on which perceptual experiences , including pure hallucinations, make present to their subjects the properties that enter into the *wide* content of those experiences.

metaphysically possible pure hallucinations do this. So this picture of what a hallucination is like is just *false*. What is going on 'in the visual field', or in other words, with the items an experience makes phenomenally present to its subject, cannot be the same in the case of a pure hallucination and a successful perception. For the latter will, and the former will not, make a perceptible scene present to its subject.

Why think that people's endorsement of the HMT might be motivated primarily by commitment to the idea that hallucinations just are cases where what is going on in the visual field is in some sense the same, and yet the world is not as it appears? Largely because, this seems to be a view which people find it extremely easy to fall into when thinking about hallucination and illusion. (In the remainder of this argument, I draw heavily on Martin (forthcoming, especially ch.1) In fact, I think the argument I give here can basically be attributed to Martin in that work.) I think that this is shown by the relative popularity of arguments from hallucination (and illusion) for the sense-datum theory, amongst analytic philosophers in the first third of the 20th century[121].

Proponents of the sense-datum theory often supported their belief in the existence of sense-data by pointing to the possibility of illusory or hallucinatory experiences[122]. These arguments usually went something like this. Firstly the sense-datum theorist would claim it was possible, for a subject *S* to have an experience 'of' some object as being *P*, where *P* is some visible property, even whilst there is no object which is *P*, that *S* actually perceives. Then it was claimed that this shows that there must be some other entity, which bears, *P*, which *S*'s experience is an experience of, and which *S* 'directly perceives'. This mysterious entity was named a sense-datum. Finally, it's then argued that if we

[121] Of course these sorts of arguments go back much further. For example, Hume is often said to give what is in essence the same sort of argument from illusion as the sense-datum theorists give in the *Enquiry Concerning Human Understanding* (see Martin, forthcoming, ch.1, pp.11-16.) Of course, this just strengthens my case for suspecting tacit commitment to the picture behind the sense-datum theory to be behind endorsement of the HMT.

[122] See, once again Martin (forthcoming, ch.1) for an overview.

directly perceive sense-data in the hallucinatory or illusory case, then we must also directly perceive them in the (fully) successful perceptual case.

Since the 1950s[123], arguments of this sort have been much criticized at the point where they move from the fact that we can undergo a hallucination which is aptly described as an experience of an object that is *P*, to the claim that in such cases there must exist some object which really is *P* that we are experiencing. It has been suggested that the sense-datum theorists simply did not understand the possibility that a phrase like 'experience of' creates a context from which existence claims are not exportable, so that one cannot legitimately infer from '*S* has an experience of a red circle', to 'there exists an object *O*, such that it is red and circular and experienced by *S*'. However, Martin argues[124], to my mind convincingly, that it is highly unlikely that the 20th century sense datum theorists simply failed to understand that some words create contexts from which existence claims aren't exportable. For example, Martin points out[125] that Russell and Moore show no tendency to make this alleged mistake in their analysis of thought, as supposed to perception. This would be odd if they were confused about the ability of mental state terms to create such contexts. If so, then, as Martin argues, it's much more likely that the sense-datum theorists were motivated by some specific conviction about the metaphysics of perception, rather than by any such mistake.

At this point, though, we can ask what sort of commitment would make the inference from '*S* had a visual experience of a red circle' to 'there is an object, *o*, which is the red circle experienced by *S*', look sound. One very salient possibility is that they are tacitly committed to the picture of what hallucinations and illusions are like, which I suggested above might underlie the apparent obviousness of the HMT. Suppose that people have a model in their minds of hallucinations and illusions as being cases where things are the

[123]Martin (forthcoming, ch1.3)).
[124] Forthcoming ch.1.
[125] Forthcoming, ch.1, p.23.

same, with respect to what's going on in the visual field as in successful perceptions, but the ordinary physical space around the subject fails to contain those items which show up in the visual field[126]. Then it's easy to see why one might be disposed to infer from the fact that someone can have an experience of, say, a red circle, which is hallucinatory, to the conclusion that there must be something red and circular which they experience[127]. Furthermore, it's not clear what other model of hallucinations and illusions could be doing the work in supporting the key controversial assumption in arguments from illusions and hallucination. If so, then we have good reason to think that, at the very least, early 20th century sense-datum theorists were (either tacitly or explicitly) operating with this model of illusions and hallucinations.

This at least suggests that such a commitment might also be at work when people find the HMT to be obviously true. After all, if this model is correct, then there really would be cases of hallucinations which perfectly mirror successful perception in terms of the qualitative nature of what they make present to their respective subjects. And we've already seen that it is part of our notion of phenomenal consciousness that two experiences will match phenomenally if there is no qualitative difference in what they make phenomenally present to their subjects. If this were the only or main grounds on which we find the HMT to just be intuitively obvious, our commitment to it ought to be at least somewhat undermined, given the inaccuracy of the model.

Of course, whether human beings (or at least analytic philosophers of perception during both the early 20th century and the present day) really do, at least tacitly, endorse such a model of hallucinations and illusions is ultimately an empirical question. I have not given nearly enough evidence here to even come closing to demonstrating that they do. Still, I

[126] Or at least, the presence of such objects in the subject's surrounding is not responsible for their experience being as it is, or not appropriately so responsible for the experience to count as a perception of all the items it seems from the inside to be a perception of.

[127] Assuming anyway, that subjects also at least tacitly believe that particular perceived objects (and not just properties) are phenomenally present in the successful perceptual case.

> think that the evidence provided for this by the case of the sense-datum theorists is at least suggestive. Furthermore, I'm hardly the first philosopher to suggest that such a model of hallucination and illusion might be somehow implicit in folk psychology.[128] Take Sellars' influential early discussion of our folk theory of sensory states, for example. It attributes to ordinary subjects something like the assumption that there is a single class of objects which fill the visual field in both the successful perceptual and the hallucinatory cases. It also attributes to ordinary subjects the belief that in cases where things in the world are other than they appear, these internal objects bear analogues of the properties which appear to be instantiated by items in the world. And Dennett (1991, pp.351-2) also makes the claim that when people are confronted with, for example, colour illusions, they have an innate tendency to think that since they experience the colour involved in the illusion, but are not perceiving it, there must be some alternative 'inner' region that is coloured in the relevant way.

Neither of these three arguments, is, of course, at all decisive, but I think that, taken together, they provide us with reasonably good reasons for preferring Phenomenal Difference views to Anti-Presence views. In the next Section, I develop my own favoured Phenomenal Difference View.

Section Two-An (Almost) Representationalist Phenomenal Difference View

My own preferred Phenomenal Difference view draws on a mixture of naive realist and representationalist ideas. The thought is that, roughly speaking, to have a phenomenally conscious experience is to successfully represent one or more objects as being a certain way. In

[128] Of course, the argument here doesn't even need a claim this strong. It's enough if the tacit model is shared by philosophers.

turn the content of this successful perception determines a perceptible scene, *S* which is thereby made phenomenally present to you by your experience. The phenomenal properties of experiences are then taken to supervene on the items which they make phenomenally present to their subjects. I call this view the 'successful representation account of phenomenal consciousness', (hereafter **SRAPC**). In this Section, I develop the view. In the next section, I consider various objections to it. In particular, the SRAPC involves the idea that phenomenal properties are, basically[129], properties of successfully[130] representing the world to be a certain way. Since it follows from this that the SRPAC is inconsistent with the HMT, the SRPAC may be thought vulnerable to various standard objections to naive realism. And since the view treats perceptual experiences as having representational content, I also need to show that various arguments given by philosophers such as Brewer and Travis, against the view that experiences have representational content do not refute the view or seriously damage its credibility.

A brief overview of what I take to be the overall argument supporting the SRAPC is needed. Firstly, I have already given reasons for preferring a Phenomenal Difference view over either a Hallucinatory Presence view or an Anti-Presence view. So the only work still to do in this Section and the next in giving a positive reason to endorse the SRAPC is to show that it is the best available Phenomenal Difference view. My attempt to do this is woven into my response to objections to my view in the next Section. Some of those objections I think genuinely do cause trouble for alternative Phenomenal Difference views on the market. If this is right, then by showing that my view avoids these objections, I do not merely show that it's tenable, but also that it's superior to extant Phenomenal Difference views. So my overall case for the SRAPC is a combination of the following. A) the case already made for adopting a Phenomenal Difference view over an Anti-Presence view or a Hallucinatory Presence view. And b) the fact that that some

[129] As is well known, since not all representation is conscious, technically, unless we find some sort of representational content possessed only by phenomenally conscious states, phenomenal properties can only be properties of 'having X representational property and also being a mental state such that...', rather than simply being representational properties.

[130] Where success requires successful perceptual contact with the world, not just veridicality.

of the objections I deal with in the next Section, whilst lacking the power to significantly damage the SRAPC, do cause significant problems for rival Phenomenal Difference views. Overall, what hopefully emerges is a view that has following advantages. It is able, unlike Hallucinatory Presence views, to get the facts right about the sort of items, whole perceptible scenes, which are present in the successful perceptual case. Unlike Anti-Presence views, it is able to do so without forcing us to deny that two experiences match phenomenally only if there's no qualitative difference in what they make phenomenally present to their subjects. And it also has the explanatory advantages which I think accrue to views which allow that all kinds of perceptual experience, successful perceptual, illusory or hallucinatory, bear representational content.

With that in mind, we can now turn to the details of how the SRAPC works. I'll first deal with how the view treats successful perceptions, and then with how it treats pure and impure hallucinations, and illusions[131].

First, successful perceptions. In this case, the idea is quite simple: ***successful perceptual experiences bear singular content, and that singular content will determine a perceptible scene constructed out of the properties and objects which enter into the singular content. That perceptible scene is then made phenomenally present to the subject of the experience***. The phenomenal properties of the experience in turn are fixed by the qualitative aspect of the perceptible scene that is present[132]. Why think that, if experiences have some kind of singular content, it will determine a perceptible scene? Well, recall what perceptible scenes are supposed to be. Their atomic parts are objects, properties, and relations, and these are glued together into perceptible scenes by the instantiation relation holding between objects and properties, and the

[131] It's not clear to me whether the SRAPC can be extended to (putatively) phenomenally conscious experiences that are not, or not obviously, perceptual, such as moods, emotions, occurrent thoughts, sensory imaginings and so forth, and I won't discuss the issue here. Obviously if it cannot be, it loses some plausibility, although how much is not clear.

[132] Though if someone wants to argue that in addition to the phenomenal properties which are fixed by what's qualitatively going on with the items which are phenomenally present, experiences also have further phenomenal properties which are fixed by which objects are present, independently of what's going on qualitatively, I would not especially object.

holding of relations between two or more of the objects. It helps here to think in terms of the structure diagrams I introduced in ch.2. Perceptible scenes can of course be represented by structure diagrams. On those structure diagrams the nodes stand for objects, the predicates at the nodes stand for the properties connected to those objects by the instantiation relation, and the relation terms on the lines represent the relations which bind the objects together into the scene. Now, any experience which bears singular representational content, will 'say' of certain object(s) that they have certain properties, and/or that they stand in certain relations to each other. So it will therefore lay down a condition which perceptible scenes can either meet or fail to meet: that they are a scene consisting in just those objects which the experience 'says' instantiate particular properties instantiating those properties, and related to each other as the experience 'says' those objects are. And in the successful perceptual case, where what the experience says about the world is correct, there will be a perceptible scene which meets this condition. So the experience's representational content will determine a unique perceptible scene[133].

As stated above, this view of the successful perceptual case entails the denial of nominalism. For since perceptible scenes as I have defined them have properties and relations as parts, any claim which entails the existence of perceptible scenes will entail that nominalism is false. Since this is a highly controversial metaphysical claim it would be nice to have a way of nominalising the view. Talk of perceptible scenes could then treated by the nominalist as an in-principle dispensable fiction.

Note first that the problem of apparent inconsistency with nominalism which occurs for the SRAPC here is not unique to the SRAPC alone. We've already seen, recall, in ch.2, that any nominalist view has to give some sort of account of the fact that experiences, even successful perceptual experiences of the same objects, and the same surface areas of those objects, can

[133] I suppose one might worry that there could be unusual perceptual experiences such that they represent two or more objects without those objects being represented as related to each other in any way, or connected by a chain of relations between objects. But if so it's not clear why we can't just hold that there are multiple scenes present in such cases.

differ phenomenally. The obvious solution to this, as I said in ch.2, if we are also committed to transparency, is to endorse the following claim. The phenomenal properties of experiences depend partly on their making objects phenomenally present to their subjects, and partly on how they are related to those objects[134]. For example, one might hold that in the successful perceptual case to have an experience with the phenomenal property associated with seeing things that are $blue_{21}$, for example, is to have to an experience *E* which meets the following two conditions. Firstly it makes an object *o* phenomenally present to one. And secondly *E* is related *R*-ly to that object *o*, where being related *R*-ly to an object fixes that the object is experienced as $blue_{21}$, rather than as having some other property. One very obvious way of cashing this out looks exactly like a nominalisation of the view of successful perceptions just sketched: take being related *R*-ly to an object to simply be *successfully* representing it as $blue_{21}$. Here successfully representing an object to be a particular way entails not just accurately representing it to be that way, as a veridical illusion might, but also counting as a perception of that object being that way. The view of successful perceptions then becomes the following. Experiences bear singular content; the objects which enter into their singular content are the items they make phenomenally present to their subjects; and it is the phenomenal presence of those objects, *and how they are represented to be,* which determine the phenomenal properties of the experience. So not only is the problem of how to nominalize things here not specific to the SRAPC, but there is an obvious solution to it[135]. Below, when I explain how the SRAPC handles illusions and pure and impure hallucinations, I

[134] Nominalists can of course accept that there are facts about how things are related, without reifying relations.

[135] An alternative which is still, I think, compatible with the argument of ch.s2-3, but which would make the nominalization look less like the SRAPC on its platonist reading, would be the following. Instead of taking experiencing an object to be *P* as a matter of *successfully* representing it to be *P* we could simply take it to be a matter of representing it to be *P*. This would not help preserve the HMT, but it would help us preserve a view on which illusions can match successful perceptions. Once a view drops the HMT, I'm not sure it gains all that much by holding that successful perceptions and illusions can match phenomenally. For in giving up the HMT, we have already dropped the basic picture behind it, on which whether an experience makes perceptual contact with any objects, or properties of objects, in itself makes no difference to the experience's phenomenal properties. Once we have given up this image, it's not clear to me why we should retain its consequences in the illusory case alone. So I don't see strong reasons for preferring the formulation in terms of representation alone to the formulation in terms of successful representation which more closely approximates the SRAPC on its platonist reading. A possible exception here is if illusions

will phrase things in terms of perceptible scenes rather than the nominalization given here. But I think it's mostly clear how to nominalize what I say about the illusory and hallucinatory cases along the lines of what I've said here about the successful perceptual case.

Note that the SRAPC is not supposed to be a full explanation of what it is that makes a perceptual state phenomenally conscious, or what it is for a perceptual state to have certain phenomenal properties. Many philosophers and psychologists think, there are unconscious perceptual states. If so, being a mental state that bears accurate singular content and is caused in the right way to count as a perception of the items which enter into that content, will not be a sufficient condition for being phenomenally conscious. Nor therefore will it be a sufficient condition for making items phenomenally present to a subject.

Instead, I conceive of the role of the SRAPC in the following way. Any account of the metaphysics of phenomenal consciousness must be able to tell us why some particular mental states are (phenomenally) conscious and others are not. And it must also be able to tell us why particular conscious states have the phenomenal properties they do. The SRAPC is supposed to help with the second task, rather than the first. In particular, it should be obvious by now that I think that the answer to the second question is that the phenomenal properties of conscious perceptual experiences are constitutively fixed by, and hence supervene on[136], the items those experiences make present to their subjects. I conceive the SRAPC as in turn giving us an account of the non-phenomenal supervenience base of facts about which items conscious experiences make phenomenally present to their subjects. In doing so, given the transitivity of supervenience, it also gives us a non-phenomenal supervenience base for facts about which phenomenal properties conscious experiences possess. Hence the SRAPC is not a full account, even of the constitutive

---

turn out to be extremely pervasive in the manner discussed in the last Section of this chapter. See the end of Section Four.

[136] I take this supervenience to be metaphysical, but we can imagine a dualist version of the SRAPC which takes it to be nomological.

grounds of phenomenal presence, but the hope is that it could be a large part of such an account[137].

So, the basic idea behind the SRAPC is that phenomenally conscious perceptual states have the properties they do in virtue of them making successful perceptual contact with objects and their properties, and so correctly representing them to have certain properties. So far, I have only applied this idea in the successful perceptual case. But I will now extend the theory to cover pure and impure hallucinations, and illusions.

First, we can take the case of pure hallucinations. The SRAPC can treat pure hallucinations in what I think is the most natural way, as experiences which represent the world as containing one or more objects and represent those objects as being certain ways[138] but are not caused in the right way to count as perceptions or misperceptions of any objects in the world. In the next Section, I will argue that this allows for a better treatment of hallucination than does the view associated with Martin (2004) that pure hallucinations are simply experiences which cannot be told apart from successful perceptions by introspection alone[139].

Note that, in itself, this not a full definition of what it is to be a pure hallucination. The reason for this is that just as we can ask what separates phenomenally conscious successful perceptions from unconscious perceptual states, so we can ask what separates pure hallucinations from early visual states which involve failed attempts at singular reference. Now we might simply carry over our response to this difficulty in the successful perceptual case, and say that a pure hallucination is an internal representational state that is in some sense a perceptual state, but that makes some

[137] The idea of using the representational content of an experience, *E* to explain particular phenomenal facts about what phenomenal properties an experience instantiates, or which items it makes phenomenally present to its subject, whilst using other features of *E* to account for what makes *E* phenomenally conscious at all, is a familiar one from the literature on representationalism. See Tye (1995, ch. 5) and Lycan (1996, ch.4).

[138] Talk of representing objects here is obviously to be read *de dicto* not *de re*!

[139] Technically speaking, Martin only commits in (2004) to this being true of a particular restricted class of pure hallucinations, those which have the same proximal causes as successful perceptions.

failed and no successful attempt at singular reference to an object. It's not clear whether or not this is right though, as we can see by considering the following case. Suppose some drug caused internal representational states which caused people to report perceiving objects in the world, and behave as if they had done so, but it turned out that those internal representational states lacked whatever property, *P*, distinguishes successful perceptions from unconscious perceptual states. It's not clear that such a case is possible, since it might be that reliably causing verbal and non-verbal behaviour similar to a successful perception is itself enough to possess *P*. If it is possible however, it seems reasonable to call the drug-induced experiences in such cases hallucinations. And they might in principle be experiences which involve no successful singular reference to objects in the world at all, and so pure hallucinations if hallucinations at all.

I shan't try and resolve this problem here, giving a full account of pure hallucinations. Instead, I will rest content with saying that pure hallucinations are experiences with the following features. Firstly they have existential but not singular content, (existentially) represent the world as containing one or more objects. Secondly they represent those objects as being certain ways but are not caused in the right way to count as perceptions or misperceptions of any objects in the world. And thirdly they meet some further functional condition which distinguishes them from the occurrent thoughts or unconscious mental states which meet the first and second criteria.

Of course, as well as the positive account of pure hallucinations which I've just given, it should by now be obvious that the SRAPC also includes a particular negative claim about pure hallucinations. According to the SRAPC, pure hallucinations do not make perceptible scenes phenomenally present to their subjects. Therefore no metaphysically possible pure hallucination can match a successful perception phenomenally, given that the QRM is also part of the SRAPC.

However, whilst I take myself to have shown in ch.s2 and 3 that it's impossible for pure hallucinations to make perceptible scenes phenomenally present to their subjects, this does not yet settle the question of whether or not they make *any items at all* phenomenally present to

their subjects. One possibility is that pure hallucinations make nothing at all phenomenally present to their subjects. On such a view, given the dependence of the phenomenal properties of successful perceptual experiences on the phenomenal presence facts about such experiences, this would mean that pure hallucinations have no phenomenal properties whatsoever. Another possible view however is that pure hallucinations make present to their subjects some of the items which enter into the existential content which they do bear. This could be the properties which enter into that content, or perhaps even the whole existential Russellian propositions which they bear as content. This account of pure hallucinations fits well with the connection which the SRAPC draws in the successful perceptual case between representing items and making them phenomenally present to a subject. On this sort of view pure hallucinations might then inherit some phenomenal properties or other from those items which they did make present to their subjects. They still would not, of course, be able to match successful perceptions given that they don't make perceptible scenes present to their subjects[140].

I think both options here have their advantages and disadvantages. But for reasons of space, I won't try to decide between the two views. Note that since the rival Phenomenal Difference view I discuss in the next Section treats pure hallucinations as having no phenomenal properties at all, at the very least the SRAPC cannot be worse off than those views here. Either the best option here, given the endorsement of the rest of the SRAPC, really is to hold that pure hallucinations possess no phenomenal properties. If so the SRAPC is no worse off, in this respect, than the rival Phenomenal Difference view which says the same thing. Or, alternatively, there is a reasonable way of working out the claim that pure hallucinations do make some items which enter into their representational content phenomenally present to their subjects. But if so, it seems doubtful that it could be an advantage for a rival Phenomenal Difference view that it held that pure hallucinations make nothing present to their subjects and lack phenomenal properties altogether.

[140] Note that all of these views on which pure hallucinations still make some items phenomenally present to their subjects require the existence of properties, and so are inconsistent with nominalism.

Surely it doesn't count in favour of a view that it affirms that counter-intuitive doctrine, rather than the doctrine that there is a mere *difference* in what pure hallucinations and successful perceptions make phenomenally present to their subjects. So whichever of the two options we choose, it's unlikely that there is some reason here to prefer a rival Phenomenal Difference view to the SRAPC

So much for pure hallucinations; what about impure hallucinations? Here I think the best tactic is to treat impure hallucinations as exactly what they seem like. Impure hallucinations are a combination of successful perception of objects[141] -and (pure) hallucination.

In particular, this, I think, is what is going on when a subject, *S* undergoes an impure hallucination, *H*. *H* is a perceptual experience of *S*'s. Since it is a perception of one or more objects it will carry singular content about those objects. At the same time, since it is partly hallucinatory, it will also presumably involve a failed attempt at singular denotation of one or more objects. I think that is best captured by treating the vehicle of content as some kind of internal sentence, containing singular terms. Some of these singular terms succeed in denoting objects in the world, and others which fail to do so, as empty names in natural languages fail to denote. This sentence will contain singular terms which denote the objects which the experience is a perception of within it, whilst also containing empty singular terms corresponding to the objects being hallucinated. For example, if I undergo an impure hallucination, *H*, of a green triangle, next to a blue circle, which I genuinely do perceive, a toy version of such a sentence, **L**, might be the following:

L: 'O_1 is green and triangular and two feet in front of me, and O_2 is blue and circular and three inches from O_1'[142]. Here 'O_1' fails to denote anything, but 'O_2' succeeds in doing so.

[141] Not necessarily fully successful perception of the properties of those objects, since there can be impure hallucinations which are also illusions

[142] This is not meant to be a realistic with respect of the actual fineness of grain of perceptual content, or how spatial relations between objects are represented within it

In turn, the experience can be assigned an existential content, by, in effect, quantifying out the names in **L**. In the next Section, I'll show that because the SRAPC posits that impure hallucinations carry existential content, it is in a better position to account for various facts about hallucinations than various rival Phenomenal Difference views which do not posit the presence of any representational content in the impurely hallucinatory case.

So much for what is going on, representationally, in the impurely hallucinatory case. What about what is going on with respect to phenomenal presence, and therefore the phenomenal properties of impure hallucinations? Here, I think that, appropriately, the correct account of the purely hallucinatory case is to combine our accounts of the successful perception and impure hallucination.

Suppose someone perceives an object, *o*, via undergoing an experience, *E*, which is also an impure hallucination, and that *E* represents *o* accurately as having property the property, *P*. Note here that the accurate singular content of the experience will still determine a perceptible scene. An experience *can* be an impure, rather than pure hallucination, because it is an illusory misperception of an object. But we can assume for the purposes of simplification, that *E* is caused in the right way to count as a perception of both *o* and *P*, not just *o* alone. Below, once I give my account of illusions, I think it will be obvious how it can be combined with the current account of impure hallucinations to treat experiences which are a combination of impure hallucination and illusion. Since *E* is caused in the right way to count as a perception of both *o* and *P* we can treat the case as involving the phenomenal presence of the perceptible scene which has *o* and *P* as parts. Meanwhile we can treat the hallucinatory element of the experience as contributing whatever pure hallucinations contribute in the purely hallucinatory case discussed above.

Return, for example, to the case where I hallucinate a green triangle, whilst perceiving a blue circle to its left, and suppose that my experience correctly represents the colour and shape of the blue circle. Then according to the SRAPC, there will be a perceptible scene phenomenally present

to me. And that scene will consist in the blue circle related to a particular colour property and a particular shape property by the instantiation relation. Meanwhile, depending on what we say about the purely hallucination case, *E*'s being a hallucination of a green triangle may contribute certain colour and shape properties[143] to the items phenomenally present to me. If so those properties will not be bound into any perceptible scene. Or it may have no consequences for what is present to me whatsoever. The crucial thing to realize here is that *whenever* someone perceives an object accurately to have any property, there will be a corresponding perceptible scene which consists in that object instantiating that property. This is because, an experience can only accurately represent a particular object as being *P* if that object is in fact *P*. And further it follows from the definition I gave of 'perceptible scene' that whenever an object actually does instantiate a perceptible property there will exist a perceptible scene consisting in that object instantiating that property. So any experience with a successful perceptual element will always determine a perceptible scene[144].

One possible consequence of this account of impure hallucinations which are also fully successful perceptions of some object that some may find a little odd is this. It implies that what is made phenomenally present to us by an impure hallucination can be 'fragmented' in the following sense. Suppose that I undergo the following impurely hallucinatory experience, *H*. *H* is a hallucination, of a blue circle, between a red triangle, and a green rectangle, against a black wall, where the triangle and the rectangle are fully successfully perceived. Then, clearly, whatever is phenomenally present to me, it's not a perceptible scene consisting of an appropriately related red triangle, blue circle, and green rectangle, since there is no blue circle which I am perceiving. At the same time though, my experience will not be equivalent, in what it makes present to me, to a fully successful perception of just the triangle and the rectangle from the same point in

[143] Or even a fragment of a Russellian proposition, if there are such things!
[144] Though of course because of the possibility of veridical illusion, there can be cases where the singular content of an experience determines a perceptible scene, but the experience doesn't make that scene phenomenally present to the subject of the experience

space. This does not follow even if we take the line that when I hallucinate an object, nothing is thereby made present to me. For, if I were just to perceive the triangle and the rectangle, I would also perceive the segment of the black wall which is apparently occluded by the blue circle when I undergo *H*. If we think of perceptible scenes on the model of areas of space which are then filled with objects and/or properties, this may cause puzzlement. For there seems to be a 'hole' in this space, corresponding to the area which is apparently occluded by the blue circle. Certainly, we cannot visualize what such a space would be like.

But in fact there is no problem here, so long as we remember that perceptible scenes are not literally regions of space, even though they contain spatial relations as constituents. Rather they are more like facts. Perceptible scenes can capture the spatial aspects of our experience, as we saw in ch.2.3, better than propositions or complex properties can, because they have as parts, objects related to each other by spatial relations. But whole perceptible scenes themselves don't literally have spatial properties. Note that it *may* be that in some of the cases where there are 'holes' of this sort, the singular content of an impure hallucination may in fact determine multiple perceptible scenes. If multiple distinct scenes which are not bound together into a single scene are indeed present, it will be because of the 'hole' that spatial relations don't bind them into a single scene. For example, in the above case it's not obvious that the subject of *H* successfully perceives any spatial relation between the triangle and the rectangle, though it's not obvious that she doesn't .

So much for the case of impure hallucinations. How does the SRAPC deal with illusions?

The basic idea is that an illusion is an experience which is caused in the right way by an object *o* to count as a perception of *o*, and which represents, through its singular content, that *o* has a certain property *P*, but which does not count as a perception of *o* as *P*. It fails to count as the latter either because *o* is not *P*, or because its being *P* isn't causally involved in bringing about the experience in the right way. But of course, this does not tell us what is going on, phenomenal

presence-wise, with illusions. There are two obvious things to decide here. The first is whether undergoing an illusory experience of a particular object instantiating properties, or some objects standing in particular relations to each other, is a way of having a perceptible scene made phenomenally present to one. Secondly, suppose we endorse a negative answer to the first question. Then we must decide whether we should hold that having an illusory experience of an object does not, by itself, make anything phenomenally present to us, or that it makes present to us something other than a perceptible scene.

On the first point, I hold that illusorily misrepresenting an object to have some property[145] is not a way of coming to make a perceptible scene phenomenally present to a subject. The reason for thinking this is obvious. Except in the unusual case of veridical illusion, there will be no such perceptible scene, since the objects in question will not in fact be as I represent them to be, and so cannot be embedded in such a scene[146].

Now showing that illusions don't make present to their subjects perceptible scenes which have the objects illusorily misperceived as parts doesn't quite show by itself that they make no perceptible scenes whatsoever phenomenally present to their subjects. Suppose we really do live in the sort of universe Tegmark describes where 'everything that can happen, does happen'. Then one option might be to take illusions to make present to their subjects perceptible scenes which perfectly fit their existential content. I reject this view however, because I suspect that it will be subject to the same sort of problems as the Quasi-Descriptivist view discussed in the previous chapter. In particular, take the case where an experience is a mixture of successful perception and illusion. I suspect that in such cases it will not be clear when the experience makes present to

[145] Or two objects to be related in some way.

[146] I suppose someone might try to hold that in cases of veridical illusion alone such a scene is made present. But building in an exception for veridical illusion alone here looks like a way of making the SRAPC kludgier and less principled. In particular it would, I think, make the relation, *R* which holds between experiences and the items they make phenomenally present to their subjects much less natural, in the Lewisian sense, than it would otherwise be.

its subject a scene containing the items that are perceived, and when it makes present a scene which fits it existential content, but lacks those items[147].

We can assume then, that, by itself, illusorily representing a particular object, *o* to have a property (or multiple properties, or two objects as related in some way, etc.), doesn't make any perceptible scene phenomenally present to one. The question that remains is, does illusorily representing an object *o* as *P* lead to an experience making some items present to its subject other than a perceptible scene? Here, as in the purely hallucinatory case there are two options, with their own advantages and disadvantages. The first is to hold that in such cases, either *o*, or *P*, or both, or maybe even all or a fragment of the experience's singular content which has *o* and *P* as parts is made present to the subject. Or alternatively, we might hold that nothing is made so present. As in the purely hallucinatory case, the first option accords with our intuition that engaging in illusory perceptual contact with an object still makes something phenomenally present to us, and leads to our experience possessing particular phenomenal properties. But endorsing it does mean that we have to introduce the following sort of distinction between phenomenal properties. Those phenomenal properties an experience possess in virtue of making a bare property *P*, or a bare object, *o*, present to its subject, will both differ from at least one phenomenal property it instantiates in virtue of making a perceptible scene which has as a part *o* linked to *P* by the instantiation relation so present. The second option meanwhile, allows us to avoid these arguably hard to understand distinctions between phenomenal properties. But it does so at the cost of committing us to the claim that illusorily representing an object to have a property doesn't in itself make anything phenomenally present to us. As in the discussion of hallucinations above, I won't try to decide this issue here, and nothing in the argument of the rest of this chapter turns on it.

[147] In the rare case of veridical illusion there will of course be a scene which contains the perceived objects *and* fits the experiences existential content.

Finally, it's worth spelling out that I handle the mixed case where an experience is both a successful perception of the properties of one or more objects[148], but also partially illusory with a similar 'mixed' account as I give of impure hallucinations. In other words the account I give is a mixture of my account of fully successful perception, and my account of illusion itself. I take it that such 'mixed' illusions make present to their subjects any perceptible scenes they are perceptions of, whilst also, through their illusory aspect, making present to their subjects whatever we decide illusions do[149].

I have now finished expounding the SRAPC, having explained what it says about fully successful perception, illusion, and pure and impure hallucination. Before moving on to explaining why I think it is superior to alternative Phenomenal Difference views, and rebutting objections to it, however, I will say a little more about its relationship to, respectively, naive realism and representationalism.

First naive realism. In one sense the SRAPC is very close to naive realism. It treats fully successful perceptions as making ordinary physical objects phenomenally present to their subjects. It bars hallucinations from making objects phenomenally present to their subjects. And holds that there is a relation, *R*, which perceptual experiences must bear to any objects they make present to their subjects, and which experiences can only bear to objects they count as perceptions of. This relation being the relation which experiences bear to items which enter into their singular content. Furthermore, since it includes the QRM, it treats the fact that perceptible scenes will be phenomenally present in the successful perceptual case, but cannot be present in the purely hallucinatory case, as entailing that successful perceptions cannot be matched by pure hallucinations.

[148] Or the relations holding between two or more objects

[149] i.e. perhaps nothing, perhaps bare objects and bare properties and relations, perhaps fragments of singular propositions.

At the same time the SRAPC is inconsistent with some views which fly under the naive realist banner. Firstly, it represents a rejection of the view that experiences don't have representational content. Secondly, it even assigns a constitutive role to the representational content of experiences in fixing the phenomenal presence facts about perceptual experiences, and hence, in turn as fixing facts about their phenomenal properties. It also represents a rejection of Martin's (2004) idea that (certain) pure hallucinations lack any substantive 'mental' properties altogether, since according to the SRAPC, hallucinations are always bearers of at least existential content. Finally, it doesn't even definitely say that in order for an experience to make any item whatsoever phenomenally present to its subject, it must count as a perception of that item. For I left open the possibility that illusorily representing something to have a *property*, or undergoing a hallucination as of an object with a property, is a way of having that property made phenomenally present to you.

As for representationalism, the SRAPC is technically not a form of representationalism, but it comes close to being one. As usually stated, and as I have defined it in ch.1, representationalism either is or at least entails the claim that two experiences do not differ in their phenomenal properties unless they differ in at least one representational property. The SRAPC however makes this claim come out false. The SRAPC treats pairs of experiences, such that one is in some respect illusory, and the other is fully successful perceptual, but which bear exactly the same singular (and existential) content, as differing in their phenomenal properties. For there will be at least one property or relation such that the successful perceptual experience makes present to its subject a perceptible scene which has that property or relation as a part, but the illusion does not. This follows from two other things. Firstly that, according to the SRAPC, it is only when we successfully perceive properties that a perceptible scene which has that property as a part is made present to us. And secondly that for any illusion, *I* there will be some property *P* which enters into *I*'s existential content but which *I* does not count as a perception of, whereas for any successful perception, *E*, *E* will count as a perception of all properties that enter into its existential

content. At the same time though, the SRAPC is certainly not far from representationalism. The phenomenal properties of experiences are treated as depending almost entirely on representational facts, since they are entirely fixed by facts about which objects an experience successfully represents and which properties it successfully represent them to have[150].

Now some representationalists, I think, could accept this failure of phenomenal properties to totally supervene on representational content with equanimity. If someone's representationalism is mainly motivated by a desire for a reductive account of phenomenal consciousness, then this failure of supervenience is not especially worrying. We can still reduce facts about which phenomenal properties a phenomenally conscious experience has to facts about how it represents the world to be, plus equally naturalistically acceptable facts about which objects and properties it counts as a perception of. Some representationalists however, do have motivations for rejecting the SRAPC here. These representationalists are committed to a certain gloss on the claim that there's a strong tie between 'how things seem' to the subject of a perceptual experience, *E*, and E's phenomenal properties. They believe the following. A) that two experiences cannot differ in their phenomenal properties, without their being a difference in 'how things seem' to those subjects, in virtue of their undergoing the respective experiences. And b) that for things to seem a certain way to the subject of an experience, in the sense relevant to a), is for their experience to be representing the world to be a certain way. From these two claims one can quickly arrive at the claim that two experiences cannot vary phenomenally without thereby varying representationally.[151] This is one place where I am prepared to simply bite the bullet, and say that all views involve us giving up on some of the apparent data. But as compensation, I note that it's unclear to me just howwell motivated the sort of argument for representationalism discussed here is. Certainly, the supposed conceptual tie between the phenomenal and 'how things seem' is, at least when things seeming a particular way is glossed

[150] Here recall, successful representation of an object entails perception of that object and successful representation of an object instantiating a property entails perceiving to instantiate that property.
[151] Byrne (2001) gives an argument of roughly this sort.

representationally, is rejected by all those who reject representationalism. This suggests the rejected claim is at least not obviously *a priori*.

Section Three-The Superiority of the SRAPC to Extant Phenomenal Difference Views

In this Section, I attempt two, related tasks. The first is to show why the SRAPC is superior to extant Phenomenal Difference views, and also, more generally, why any adequate Phenomenal Difference view ought still to allow for the existence of representational content across all perceptual cases. Secondly, I show how the SRAPC can avoid standard representationalist objections to naive realism and standard naive realist objections to representationalism. These two tasks are related. Many of the objections which have caused problems for previous forms of naive realism turn out to only be troubling for views which deny that any perceptual experiences bear representational content[152]. Showing that the SRAPC avoids the problems generated by these objections to naive realism not only shows that it is not itself inadequate, but also gives us further prima facie reason to prefer the SRAPC to other Phenomenal Difference views insofar as these are troubled by such arguments.

The plan of attack is as follows. Firstly, I show that other extant Phenomenal Difference views, unlike the SRAPC, deny that perceptual experiences of any kind ever bear representational content. I show how to transform an objection of Burge's to the claim that 'there is never any specific perceptual-state kind in common'[153] between a perception and a (pure) hallucination, into an argument for the claim that perceptual experiences of all types do bear content. This

[152] Or at least deny that pure hallucinations and illusions can bear exactly the same existential content as successful perceptions.
[153] (2005, p.2).

refutes extant Phenomenal Difference views which deny that perceptual experiences ever bear content, and thereby shows the SRAPC to be preferable to those views.

I then go on to discuss whether the SRAPC is vulnerable to objections against naive realism already given in the literature, and show that it is not. In most cases it's precisely because the SRAPC treats hallucinations as content-bearers that it's not so vulnerable. When I think these objections, whilst they don't cause trouble for the SRAPC, may cause trouble for rival Phenomenal Difference views, I will flag this. Insofar as such views look like they might be vulnerable to such objections, this provides further reason for thinking that the SRAPC is superior to other extant Phenomenal Difference views. Note, however, that, unlike in the case of my Burge derived argument, I'm not committing to the claim that these objections show the rival Phenomenal Difference views to be false. The debate around most of the objections is quite complex. The point is just that the SRAPC clearly faces no trouble from objections which *may* cause trouble for those views.

So, on to other extant Phenomenal Difference views. Recall that in ch.1.5 I said that the clearest example of a Phenomenal Difference view was the view expounded by Fish in his (2009) book, but that Brewer and Campbell may also hold Phenomenal Difference views. Martin, recall, meanwhile seems in different places, to say inconsistent things about what phenomenal properties themselves are, one of which suggests a Phenomenal Difference view, though the other suggests that his naive realism is best interpreted as an Anti-Presence view. What matters here is that all four of these others can at least be interpreted as endorsing some kind of Phenomenal Difference view distinct from the SRAPC. Hence if we are to know whether or not to endorse the SRAPC, we must see how it does against these possible rival Phenomenal Difference views whether or not they are actually the views of any particular philosopher in the literature.

Fish holds that when we undergo successful perceptual experiences, those experiences bear a non-representational relation to facts in the world[154], and this determines their phenomenal properties, and also presumably what they make phenomenally present to subjects. Fish then, gives the following account of hallucinations. Hallucinations, as a class, are 'mental events' such that they lack phenomenal properties, but produce the same 'cognitive effects' on subsequent beliefs and action, as would some successful perception which occurred in a rational subject with the same background beliefs as the hallucinator. A hallucination, *H* of an object, *o* that is *P*, is then supposed to be a mental event which has the same cognitive effects a successful perception of some *o* that is *P* would in a rational subject who shared the hallucinator's background beliefs.

Martin, Campbell and Brewer all also reject the presence of representational content, even in the successful perceptual case. Martin tries (2004, 2006), to give an account of hallucinations in terms of their being indiscriminable by introspection from successful perceptions, rather than in terms of treating them as content-bearing states. Further, he is clearly committed to the rejection even of the weaker view that they bear representational content at all, at least in the case of pure hallucinations which share all the narrow properties of successful perceptions. For in his (2004) he argues that those who are committed to his particular form of naïve realism must endorse the view that a certain class of hallucinations lack any substantive mental properties[155] whatsoever. Brewer meanwhile, explicitly denies that any perceptual states, whether successful perceptual, illusory or hallucinatory, bear representational content. He accepts something like Martin's account of hallucination[156]. Campbell's position on whether there is content in either the successful perceptual or the purely hallucinatory case is hard to determine, but he *seems*, in his (2002), to be opposed to the idea that successful perceptions bear representational content. He

[154] Or perhaps it is the subject of the experience which bears this relation. I am not confident in attributing to Fish either the view that it is experiences themselves which bear the naive realist relation to particular objects, or that it is subjects who do so.

[155] This is not the place for discussing exactly what a substantive mental property is supposed to be. Since none of my own views are stated in a way which employs this notion, problems with it are not problems for me.

[156] See Brewer (2011, p.109.)

may, though, just reject the idea that facts about the singular content of perceptual experiences help constitutively fix their phenomenal properties. Overall, the point is that all of these philosophers-with the possible exception of Campbell, seem to be committed to opposing the view that our perceptual experiences bear representational content. I shall now argue that this opposition is a mistake.

Burge, in his (2005), argues that the claim of Martin (later taken up by Brewer and Fish) that hallucinations (or at least, a certain class of them) do not bear representational content, or possess any other substantial mental properties, is inconsistent with known results in vision science. I think that Burge is right about this. I will now argue that if Burge is right and right for the reasons he gives, then we have a good argument for the claim that all kinds of perceptual experiences, successful perceptions, illusions and hallucinations, bear representational content.

According to Burge, a central strand in research on visual perception in modern psychology has been work which assumes that vision is a process which takes us from mental states which carry information about various features of the 2-d retinal image[157], to representations of a 3-dimensional world in front of us. Current work assumes that this involves many transitions between mental states, starting with the first mental states bearing information about the retinal image. These transitions are governed by what are called 'biases'. Here 'biases' are rules governing the transitions from one state to another. Taken together, they work (albeit imperfectly) to produce the representation of the 3-dimensional scene in front of the subject most likely to be accurate, given the nature of the 2-d, retinal image. These biasing principles make reference to the representational properties of such mental states, in that they tell us how the content of later states in the process depend on the content of earlier states.

[157] Burge does not endorse the claim that they 'represent' that image, but his reason for being unwilling to say this are somewhat idiosyncratic.

Burge argues that work within this tradition in psychology has given us good reason to hold that there are substantive mental properties in common between successful perceptions, and certain hallucinations. In particular, Burge argues that the reasonably empirically successful theories produced in this tradition, all assume the following. If you hold constant a perceiver's background psychological states, then an identical proximal stimulus at the surface of the body[158] will produce the same internal representational states. This assumption according to Burge, plays a role in how psychologists collect data on which to test and build theories. In particular, in certain experiments, impurely hallucinatory conscious perceptual experiences are deliberately produced, via the production of particular input at the retina. The results of such experiments are then used to test claims about biasing principles in the following way. The hallucinations are treated as content-bearing states, with what is hallucinated telling the experimenter (roughly) what the content of the hallucination is. Interpreted this way the experiments give us information about how the proximal stimulus at the retina used in the experiment is converted into an internal representation of a 3-dimensional world of objects in these *particular* cases. This information can then be checked for what *general* stories about how the brain gets from the retinal image to 3-d representations of the world, *and in particular what accounts of biases*, are consistent with it. The accounts of the biases which are being tested against the experimental evidence in this way are accounts which assume that the same biases operate in the successful perceptual case, as well as the illusory and hallucinatory cases. Burge takes this sort of work to somehow presuppose the 'same proximal stimulus, same existential content of eventual perceptual state' principle, which is supposed to apply to all kinds of perceptual experience, whether successful perceptual, hallucinatory or illusory. Since such a principle entails that pure hallucinations possess existential content it is obviously inconsistent with the claim that pure hallucinations have no substantive mental properties.

[158] Not just the retina alone, because of multimodal interactions.

I won't discuss whether Burge is right that such work somehow presupposes that such a principle is correct. What I am instead interested in establishing here is the claim that all perceptual experiences, including pure hallucinations bear representational content. And if Burge is right in his account of the vision science, as I think he is, then we can indeed consider this claim established. If facts about the content of a particular hallucination, *H*, can be evidence for general principles about which sorts of proximal stimuli produce which sorts of content bearing perceptual states, then clearly *H* must bear content. And if the general principles hold across the successful perceptual, hallucinatory, and illusory cases, then there must be some univocal sense in which experiences of all these kinds possess representational content. This shows that Phenomenal Difference views which deny that perceptual experiences of all kinds have representational content, or deny that pure hallucinations have any 'substantive' mental properties, are in trouble. The SRAPC meanwhile, does just fine, since it posits that all perceptual experiences, including pure hallucinations bear at least existential content.

Admittedly, whilst the presence of representational content in the purely hallucinatory case serves to refute the views of Martin, Fish, and Brewer, it certainly does not show that the SRAPC must be correct if any Phenomenal Difference view is. It's possible to believe that perceptual experiences bear representational content, and still not claim that for an experience to make an item phenomenally present to its subject partially or wholly consists in it having representational content into which that item enters. It's coherent to suppose instead, that which items enter into the representational content of a successful perception is not connected either constitutively or causally to which items it makes present to its subject. So there are certainly coherent Phenomenal Difference views of the following sort which are consistent with vision science cited by Burge. Experiences make the perceptible scenes they count as perceptions of phenomenally present to their subjects. And those scenes, as a matter of fact, sometimes or always mirror the propositions which are the singular content of such perceptual states. But there is no deeper connection here between content-facts and phenomenal presence-facts.

Now, I don't actually need to refute such a view here, since all I am arguing for in this chapter is that the SRAPC is superior to extant Phenomenal Difference views, and not that it is the only view compatible with the discoveries of vision science cited by Burge, or that it is better than any possible alternative Phenomenal Difference view. Nonetheless, I think there are reasons to endorse some kind of constitutive connection between the phenomenal and the representational.

Firstly, in the successful perceptual case at least, there seems to be a very strong correlation between changes in the representational content of an experience and changes in the items it makes phenomenally present to its subject. One thing, I think, that has been shown by the literature on representationalism is that there are no *clear and undisputed* cases where an experience makes something phenomenally present to its subject, but clearly does not have representational content into which that item enters[159]. Of course, a lack of counter-examples here doesn't *prove* that we have covariance here. A lack of clear counter-examples on its own might just indicate that it's hard to for us to tell whether what items a successful perceptions makes phenomenally present to its subject co-varies with how that experience represents the world to be. But I think that the following gives us further evidence that there really is co-variance here. Typically, a careful description of what someone sees when undergoing a successful perception, can give us information about both the representational content of an experience, and what it makes phenomenally present to its subject. Or at least, given a background assumption that we see things by having experiences which represent them, this seems to be the case. This assumption is not especially controversial in the current dialectical context given that I've already argued that experiences are content-bearers. Of course, co-variance doesn't entail the claim made by the SRAPC that the facts about which items an experience, *E*, successfully represents play a *constitutive* role in fixing the facts about which items *E* makes phenomenally present to its subject. But if there is no constitutive connection here, then some other

[159] Though there are of course some alleged but disputed counter-examples to representationalism.

explanation for the co-variance must be given, and it is not obvious what that would be. If we reject causal dependence here on physicalist grounds, it's hard to see what other explanation springs to mind.

Overall then, the view that experiences don't bear content at all is in deep trouble for the reasons Burge cites, which is itself enough to cast severe doubt on all of the rival Phenomenal Difference views to the SRAPC extant in the literature which I know of. And there are also some reasons for thinking there's a constitutive connection between the representational facts about experiences and the facts about which items they make present to their subjects. I turn now to potential objections to the SRAPC. Firstly I will discuss whether it is vulnerable to various objections that have been made against naive realism. Then I shall examine some naive realist arguments against views on which perceptual experiences bear representational content.

A first, and very simple objection can be put like this. According to the SRAPC, if I undergo a purely or impurely hallucinatory, or illusory experience, as of an object *o* that is *P*, then my experience differs in its phenomenal properties from any fully successful perception of an object *o* as *P*. However, if my experience has exactly the same content as such a successful perception, then it will seem to me, on the inside, to have the same phenomenal properties as such a successful perception, and there will seem on the inside to be no qualitative difference in what it makes phenomenally present to me. Furthermore, illusions at least, if perhaps not hallucinations, are reasonably common. So the view implausibly entails that we are often mistaken about the nature of our own experiences.

Here, I have no particularly interesting response to make. Instead, I will simply record that I am prepared to bite the bullet on this objection, and hold that, in cases of hallucination and illusion we can be and usually are mistaken about the phenomenal properties of our experiences.

We can note though that it's hardly surprising that we are prone to such mistakes, given the assumption that the SRAPC is true. According to the SRAPC, the phenomenal properties of an experience are determined partly by the items it makes phenomenally present to its subject, and which items it makes phenomenally present to its subject, is in turn determined by which items it successfully perceptually represents. This suggests that we should not be too surprised by error in the case where an experience represents the world as containing a certain sort of perceptible scene, but the world fails to do so[160]. For it seems as though the access of any introspective belief-forming mechanisms as to which perceptible scene an experience is hooked up to, will in some sense go through the content of that experience. If so errors in how that experience represents the world to be will lead to errors in subsequent introspectively formed beliefs about what items it makes present to the subject or what phenomenal properties it has. More precisely the point can be put like this. The proponent of the SRAPC has to tell some story about how we know roughly what is phenomenally present to us in the successful perceptual case. For we can clearly tell *roughly* what shapes and colours (and spatial relations, and so on) are relevant to what is phenomenally present to us by introspection, even if we cannot tell by introspection alone whether what is present to us are objects, properties or tropes. The easiest way to explain this is to say that we simply read information about what perceptible scene is present to us off of how our experience represents the world to be[161]; this is the famous 'displaced perception' theory of introspection introduced by Dretske[162]. If we adopt the displaced perception theory we will get predictable error in the illusory or hallucinatory cases. For in those cases too, we will form our beliefs about what is present to us by reading this off from (roughly) how our experiences represents the world to be. Given, however, that what is present in the illusory or hallucinatory case differs from what is present in the successful perceptual case with the same existential content this will inevitably lead to error in the latter two cases. And presumably error about what

[160] Or at least the experience fails to make the right kind of contact with such a scene, in the case of veridical illusions and hallucinations.
[161] Though we are presumably not explicitly aware that this what we are doing.)
[162] (1995, ch.2)

is phenomenally present to us will be accompanied by error about the phenomenal properties of the illusion or hallucination in question.

I turn now to more sophisticated objections which have been pressed in the literature against naïve realism, to see if these cause trouble for the SRAPC.

We've already seen that the SRAPC, since it posits that phenomenally conscious perceptual experiences are always representational content bearing states, avoids Burge's worries about previous naïve realist views being incompatible with well-established theories in vision science. Many other objections to previous versions of naïve realism, also turn on objections to Martin's claim that hallucinations[163] lack any positive mental properties. For example, Siegel (2004) has argued that Martin's account of hallucinations as states which are introspectively indiscriminable from successful perceptions, is inadequate in the case of animals and other cognizers who are not sophisticated enough to entertain higher order thoughts about their own mental states. These animals, Siegel argues can still undergo hallucinations, and hallucinations of particular things. But Martin's view cannot account for this, according to Siegel. For in the case of such cognizers, any mental state will trivially be indiscriminable, for them, from a successful perception. For without the capacity for higher-order thoughts, they will not be able to judge that a particular experience *E* differs from some possible successful perception *E*'. A variety of similar objections have been pressed against views which claim that hallucinations lack substantive, non-relational, mental properties, and then go on to try and define what it is for something to be a hallucination in terms of it in some sense sharing the effects of some possible successful perception on future beliefs and perceptual judgements. Indeed, Fish's (2009, ch.4) theory of hallucination, which we met briefly above, is explicitly designed to be an account of this type which avoids this sort of objection.

[163] Or at least, pure hallucinations which have the same proximal causes as actual or possible successful perceptions

It's clear though that the SRAPC faces no trouble here, or at least, no more trouble than standard forms of representationalism. For the SRAPC does not define what it is for an experience to be a hallucination in terms of indiscriminability. Rather it defines hallucinations as states which bears the same functional role[164] as successful perceptions, and bear existential content, but have an empty singular element, due to a failed attempt at reference to an object. Nothing about this implies that anything indiscriminable by introspection from a hallucination is itself a hallucination, so the problem raised for Martin's account by the existence of perceptual experiences in unsophisticated cognizers does not arise. Of course, Martin has his own response to this sort of objection, which for reasons of space I won't discuss here. Suffice to say that if the reader thinks Martin's response is unsuccessful, this counts as another point in favour of the SRAPC as against Martin's view (and Brewer's) since he seems to endorse Martin's account of hallucination). But I am not taking a stand on whether this is so, my main point is just that there is clearly no trouble for the SRAPC lurking here.

Other objections to naïve realism, or views somewhat like it in the literature, also seem to turn on negative claims made about what hallucinations or illusions are not, which have been endorsed by some previous naïve realists, but which are not entailed by (or even consistent with) the SRAPC. For example, Byrne and Logue (2008, sc.8) direct some of their fire against the claim that illusions cannot share substantive mental properties with successful perceptions. The SRAPC certainly does not entail this view. Instead it's part of the SRAPC that an illusion and a fully successful perception can bear exactly the same existential and even singular content. Byrne and Logue also argue (2008, sc.9) against Martin's view that (at least some) pure hallucinations lack

[164] Admittedly, we saw in the previous Section that the SRAPC doesn't offer a full description of just what relevant functional properties make something a hallucination. For we don't yet have an account of what distinguishes conscious from unconscious representings, but merely the suggestion that such an account will somehow involve playing a particular functional role. (Though it is not as though there has been no work on what this functional role might be; classic early work on this by representationalists includes Tye (1996, ch.5.2) and Carruthers (1996, ch.s6-7)). So there may still be some objections open to the SRAPC on the basis of its failure to specify the relevant functional/cognitive role which it says is *partly* constitutive of something's being a phenomenally conscious perceptual experience. But at least it is not directly open to objections directed against the claim that occupying the specific functional role of being indistinguishable from a successful perception is all of what it is for something to be a conscious perceptual experience.

any substantive mental properties. They argue that if successful perceptions bear representational content, so do pure hallucinations. For similar reasons the SRAPC is obviously not vulnerable to this argument.

It is in fact quite hard to find direct objections to the view, entailed by the SRAPC, that pure hallucinations cannot match successful perceptions, as supposed to objections to further negative claims made about pure hallucinations by previous naive realists. It seems that, if we look for objections that don't turn directly on these features of previous naïve realist views, we are left largely with the claim that if a view entails the denial of the HMT, then it is automatically highly counterintuitive[165]. But, as we've seen throughout this Thesis, any view here which is not a Hallucinatory Availability view, will have to give up one claim or another which, at first glance looks very plausible, whether it is the HMT, or the QRM[166]. I have already given my-somewhat tentative-reasons for rejecting the HMT, rather than the latter claim above in Section Two, so I won't repeat them. Rather, I will now turn in the next Section, to arguments against the view that experiences have content. But the important point is that the mere, though correct, claim that denying the HMT is counter-intuitive doesn't alone give us decisive reasons to reject views which drop the HMT, as the SRAPC does. One must also show that there is a view that fits the other data we have about phenomenally conscious perceptual states and does not have equally (or more) objectionable consequences.

Having shown that the SRAPC is not vulnerable to standard objections to naive realism, I shall now defend it from some objections made in recent naive realist work to the very idea that perceptual experiences have representational content. My tentative defence of the SRAPC will then rest. Note that for reasons of space, I will not discuss Martin's famous (2004) 'blocking' argument, although the SRAPC clearly does fall within the range of views which that argument targets.

[165] Tye (2002) is one example of a philosopher who argues in this way.
[166] Or transparency or extended transparency.

Firstly, let's look at the case against the view that experiences bear representational content that's given by Charles Travis in 'The Silence of the Senses'. The original version of this article is Travis (2003), but I focus on the revised version which forms ch.1 of Travis (2013).

The basic idea behind Travis' argument is clear enough. Firstly he claims that those who posit that perceptual experiences bear representational content ought to hold that the subjects of those experiences have at least a rough grasp on how those experiences represent the world to be. Secondly he argues that representationalism fails this test. Travis' argument that it fails this test, in turn, depends crucially on the claim that there is no sense of 'looks' such that talk of how things look to the subject of an experience reports a fact about how that experience represents the world to be. Hence it follows that subjects cannot come to know the representational content of their (visual) perceptual experiences by reading it off facts about how things look to them.

Travis argues that there are only two notions of 'looks' available. One is, roughly, a comparative notion, where to say that someone looks *P* is to say that they look as things which really are *P* do. For example, to say that someone 'looks tired', in this sense, is according to Travis, to say that they look as tired people typically do. The second notion, on the other hand takes looks talk to be talk about what some portion of visual evidence reports, as in when someone says that 'it look as though Pia will sink the putt' or 'that painting looks to be a Vermeer'[167]. An example of this particular sort of use of the second sense of looks talk given by Travis, is of someone saying that a lemon-shaped object looks to be a lemon. Travis goes on to argue that on neither of these senses of 'looks' does talk about how things look to a subject indicate a fact about how the subject's current visual experience represents things to be. In Travis' terms representational content is not 'looks indexed'.

What reason does Travis give for thinking that neither sense of looks 'indexes' the representational content? The problem is supposed to be that it is highly implausible, on either

[167] Examples adapted from Travis' own.

sense of looks to think that if some object, *o* looks *P* to the subject of an experience *E*, then things are the way *E* represents them to be only if things are the way they look, or in other words, if *o* is *P*.

In the case of the first, comparative sense of 'looks' the problem is supposed to be that holding that if *o* looks *P*, then *o* is represented to be *P* then we will assign to our all our experiences impossible contents. The argument for this is that if an object looks like things that are *P*, it must also look like things that look like they are *P*, but are not. To use Travis' example, if something looks lemony, in the comparative sense, it also looks like a cleverly made replica of a lemon. And even if something looks red, in the comparative sense, it also looks, in the comparative sense, like a blue thing which is currently reflecting a combination of light waves of the sort usually reflected by red things. Obviously, however, our perceptual experiences do not all have impossible contents!

In the case of the epistemic hedging sense of looks, meanwhile, the complaint is as follows. If an object, *o* looks to me to be *P* in the epistemic sense, this cannot be what constitutes the experience's having a certain content for the following reason. For things to look a certain way to me (in the epistemic sense), depends on my background evidence in a way in which, plausibly, the representational content of perception, if there was any such thing, would not. For example, in a world with very few lemons, but many lemon shaped bars of soap, lemons will not look like lemons, in the epistemic sense, because seeing them will not usually provide someone with good evidence that there is a lemon about.

Travis goes on to argue in sc.6 of (2013, ch.1), that if the subjects of perceptual experiences cannot read how their experiences represent the world to be off how things look to them (in either sense of looks), then there is no other way for them to be able to determine what the representational content of their experiences is. I find this part of Travis' argument difficult to follow, but it largely seems to rest on the question 'if not looks, then what'? In other words, the

main strand in this last part of his argument seems to be the claim that if we cannot come to know how our experiences represent the world to be by attending to how things look to us, then there is no obvious other way that we can do so.

Now, there are several very obvious questions about Travis' reasoning here. For example, even before we ask what the rationale behind it is, it's not clear just what Travis' requirement that subjects be able to tell how their perceptual experiences represent the world to be amounts to. Does it mean that they must literally believe, under a guise/mode of presentation involving the concept REPRESENTATIONAL CONTENT, that their experiences have a certain content? Or just that they must be able to know statements about their experiences which a priori entail, or metaphysically necessitate that those experiences have a certain content? Or even just that they can know things about their experiences from which, in the actual world, the content of the experiences can be read-off? Travis seems to lean towards the first and strongest requirement here, given that he says things like this: 'that we are represented to in experience is meant to be...something we can tell is happening' (2013, p.52). But he does not explicitly clarify the issues, largely resting content with the vague formulation that if experiences have representational content, they should represent 'to the subject'. Similar issues arise around how to interpret the proposal, which Travis makes on behalf of those who believe the perceptual experiences have representational content, that that content is 'looks-indexed'; it's not clear whether this is supposed to involve a metaphysically necessary, or merely a contingent connection. Secondly, it's not immediately clear exactly why Travis thinks that the two senses of 'looks' which he discusses are the only senses[168].

It is tempting, however, rather than worrying about any of this, to simply dismiss the relevance of Travis' argument to the SRAPC in the following way. Travis does not directly argue for the claim that those who believe that experiences have representational content must hold that the

[168] A point also made by Byrne (2009).

subjects of such perceptual experiences are able to tell how their experiences represent the world to be (however we define 'able to tell'). Instead, he takes this claim to be simply another part of the views he is criticizing, along with the claim that experiences have perceptual content itself. However, in expounding the SRAPC, *I* have not directly claimed that the subjects of perceptual experiences can tell how their experiences represent the world to be. I've simply claimed that perceptual experiences have representational content, and assigned this content a role both in fixing the items which our experiences make phenomenally present to us[169]. But I have not directly made the claim that when we undergo perceptual experiences, we can tell how our experiences represent the world to be. Because of this, it may seem that, however threatening Travis argument may be to those accounts of perceptual experience which do make this claim-Travis cites many philosophers making it in (2013, ch.1)-it simply has no bite against the SRAPC, and we can therefore simply move on.

Alas, things are not quite so simple. Firstly, my argument in ch.2.3 relied on the claim that it's just obvious that experiences have accuracy conditions, as well as on the invocation of intuitions about the relationship between the accuracy conditions of experiences and their phenomenal specification conditions. Now, the notion of 'accuracy conditions' and 'representational content' may not be *exactly* synonymous, or *a priori* equivalent. The former is usually introduced via its connection to the ability of thoughts/belief (and perhaps perceptual states) to be 'about', things, in a perfectly normal sense of the English word 'about'. Whereas the latter is usually introduced via a notion of the possibility of perceptual error, plus the fact that things could have been the way that they are hallucinated or illusorily misperceived to be, in such cases of error. But it's clear I think that the two notions are closely akin. Travis is therefore hardly likely to accept that we can

[169] I did also tentatively endorse the displaced perception theory of introspection on which we read-off what is phenomenally present to us from the content of our experiences, when responding to one objection above. But I do not think that the displaced perception theory is a core part of the SRAPC itself, or that the argument I used the displaced perception theory to respond to is all that threatening even if we drop that response. So we don't need to worry about whether the displaced perception theory implies that subjects have some kind of access to the content of their experiences.

just know roughly what the accuracy conditions of our experiences are, even if the latter are not looks-indexed. And he is likely to run the same arguments against looks, in either of his two senses, being able to 'index' accuracy conditions, as he ran against representational content being looks indexed. Since I've relied on philosophers at least having some kind of grasp of the accuracy conditions of their experiences in ch.2.3, Travis argument does therefore still present a challenge to me after all.

Furthermore, it seems that the methodology of the programme in vision science which I described in the previous Section, does presuppose that what ordinary subject's say is some kind of guide to the representational content of their experience. For as we noted there, vision scientists take subjects reports about their conscious perceptual experiences to be evidence about the representational content of those experiences. If Travis has somehow shown, via his argument that representational content is not looks indexed, that there is no way for what ordinary subjects say about their experiences to track their representational content, he has thereby undermined the presuppositions underlying this research programme. If so he has also undermined the claim that its apparent success provides us with good reason to think experiences bear representational content. Although given that the research programme is well established and empirically successful, we should be extremely reluctant to concede to Travis any claims that are inconsistent with its underlying presuppositions.

Luckily, however, there is another quick and easy dismissal of Travis' argument available. Nothing I've said about accuracy conditions in ch.2, requires that ordinary subjects have a concept of accuracy conditions or that there are any synonyms or near synonyms of 'accuracy conditions' in English. All I relied on in ch.2 is that philosophers, once the term 'accuracy conditions' has been introduced as a piece of philosophical terminology, are capable of telling roughly what accuracy conditions certain hypothetical perceptual experiences would have. This means that it doesn't matter if Travis being right about looks would show that there are no sentences in ordinary

English that are equivalent to claims about accuracy conditions. At most what is needed is just that there are facts about experiences from which philosophers who know what an accuracy condition is, can read off what their accuracy conditions are. A similar point applies in the case of the justification vision scientists have for taking certain things that experimental subjects say to be evidence about the representational content of their perceptual experience. On the one hand, we've seen already that for something a test subject says to actually carry information about the content of their experience, it's not necessary that the subject of the experience has the concept REPRESENTATIONAL CONTENT. Nor is it necessary that what they said is *a priori* equivalent to a claim about the representational content of their experiences. All that's necessary is that there should be a way to reliably infer, in the actual world, from the test subjects statement to some fact about the content of her experience. Meanwhile the *justification* which psychologists working on vision have for taking what subjects say to reflect the content of their experiences, is that if they do so, they are able to construct theories which explain much of their experimental data and make novel predictions which are later verified. The psychologists don't need to be antecedently justified, qua competent speakers of English, or some other natural language, that the relevant statements of the experimental subjects are equivalent to or entail claims about the representational content of their experiences.

What this shows, I think, is that Travis has managed to threaten the argument of ch.2, or the presuppositions of current vision science, only if, in showing that representational content is not 'looks indexed' in his sense, he has also shown the following. Travis must have shown that there are no introspectively accessible facts about our perceptual experiences from which, as a matter of fact, we can read off their representational content or accuracy conditions. Otherwise, it will be plausible that philosophers can read the accuracy conditions of an experience off these sorts of facts about it, when making judgments about hypothetical cases of the kind I relied on in ch.3. And, provided that ordinary subjects sometimes state facts of the relevant kind about their experiences under questioning by experimenters, it will also be plausible that psychologists can

pick up on facts about the representational content of experiences by asking the subjects of those experiences about them. Once we see this though, we can see that there is a very obvious set of facts about experiences from which we might be able to read off their content, but which are not obviously facts about how things look to the subject of such experiences. These are facts, in the successful perceptual case, about what a subject sees, in the illusory case facts about how they would be seeing objects to be were the experience veridical, and in the hallucinatory case, about what they hallucinate. Plausibly, if I see an object, *o* to be *P*, then my experience *E* is accurate iff *o* is *P*, and *E* represents *o* to be *P*. Likewise if I have an illusory experience through which I would be seeing *o* to be *P*, if my experience of *o* was successful perceptual rather than illusory. And if I hallucinate that there is an object that is *P* in front of me plausibly my experience is accurate only if there an object that is *P* located in front of me, and will carry existential content which says that there is an object that is *P* in front of me. Furthermore, it's hard to see how any claims about the semantics of 'looks' in English could cast doubt on these claims. At the very least Travis does not show how they do. Nor does he show even how claims about the metaphysics of looks themselves, rather than the semantics of 'looks' could do so. Of course, statements about what we see, or how we illusorily experience an object to be, or what we are hallucinating aren't equivalent to claims about the representational content or the accuracy conditions of an experience. But this is ok, so long as we can still reliably infer the latter from the former.

Of course, claims of the form '*S* saw that *o* was *P*', '*S* had an illusory experience of *o* as *P*' or '*S* hallucinated an *o* that was *P*', might, in certain circumstances, *entail* claims about how things looked to that subject[170]. However no trouble automatically follows from this, for the view that we can read off the content and accuracy conditions of experiences from such claims about them, even if we accept what Travis says about looks.

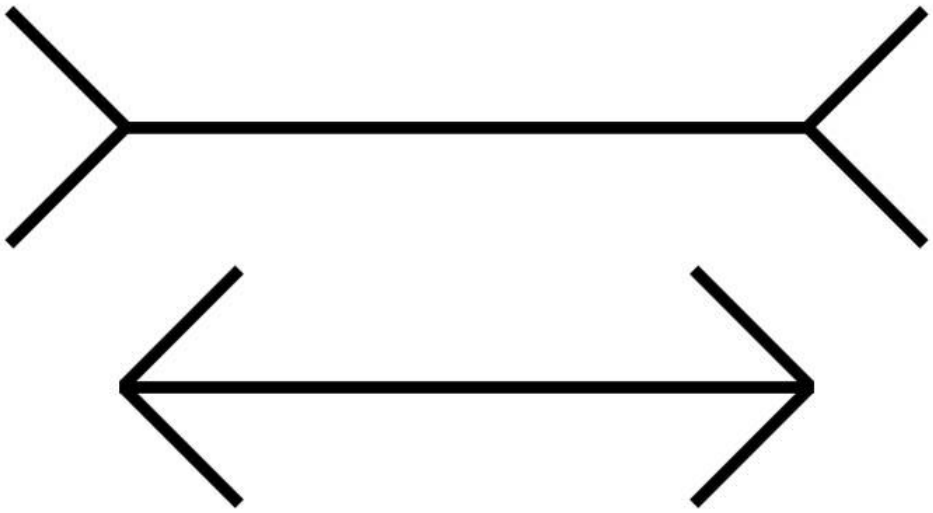
[170] That there are readings of '*S* saw that *o* was *P*' on which it entails a clam about how *o* look to *S*, if there are any such, doesn't imply that there are not also readings on which it entails no such thing.

So Travis' argument is unpersuasive. Firstly, Travis doesn't give us any reason to endorse his key claim that if perceptual experiences bear representational content, then the subject of such experiences will have at least a rough and ready grasp of that content. Secondly, insofar as my argument for the SRAPC is committed to any claim about ordinary subject's grasp of the content of their experiences, it is only the very weak claim that, as a matter of fact, statements about what we see, illusorily misperceive or hallucinate are a guide to the content of our experiences. Travis' claims about looks, even if true, don't seem to give us reason to reject this weak claim. I will now therefore move on to a second set of arguments against the view that perceptual experiences bear representational content, which are found in Brewer (2011).

Brewer gives three arguments against the view that perceptual experiences bear representational content in (2011). The first is an argument that certain illusions, such as the experiences of the Muller-Lyer line, would have to have impossible content if they had any representational content at all. Brewer argues that this clashes with a principle linking conceivability and possibility which he thinks is plausible. Secondly, he makes the following argument, translated into my terms. Suppose we take experiences to make phenomenally present to their subjects whatever objects they bear singular content about. Then we will not be able to explain the (alleged) fact that a certain level of misperception of an object's properties is incompatible with that object's being made phenomenally present to you via the experience through which you so misperceive it[171]. Finally, Brewer gives a third argument from 'the generality of predication', however I am not going to attempt to describe or respond to this argument, partly for space reasons and partly because I find it very unclear what the argument is supposed to be.

Brewer's first argument is, as I said, involves the famous Müller-Lyer illusion:

[171] Strictly speaking, this argument isn't designed to show the very strong claim that experiences lack representational content altogether. Rather it's designed to establish the weaker claim that, translated from Brewer's terms into mine, the representational content of experiences can't play a constitutive role in fixing which items they make phenomenally present to their subjects. But this is also a threat to the SRAPC.



Here the upper line looks longer than the lower to normal perceiver, but there is in fact no difference in length between the two lines. Brewer argues that those who hold that experiences have representational content must assign an impossible content to experiences of the Muller-Lyer illusion. He then argues that this is problematic, on the following grounds:

> Perhaps proponents of (CV) may simply accept this result with equanimity. It does seem to be in serious tension, though, with relatively widely accepted views about the relation between conceivability and possibility. Nobody thinks that distinct perceptions, even perceptions of a single subject over a short period of time, need conjoin consistently or possibly. But the idea that what we can conceive is a guide to what is possible plausibly rests, at least in connection with the perceptible macroscopic physical world, on a combination of the following two principles. First, that what we can conceive in connection with the macroscopic physical world is constrained by what we may experience in a single perception: very crudely, perceptual experience provides the raw materials of genuine conception in this domain and also controls the ways in which these may legitimately be employed. Second, that what we may experience in a single perception is at least a guide to what is possible: if we really are able to take something in in a single perception, then this must at least be a possible state of the world. Given these two principles, it follows that legitimate conception of this kind will be, as many believe that it is, a guide to what is possible with respect to the perceptible macroscopic physical world. (Brewer, 2011, p.67).

Brewer's reason for thinking that when someone undergoes the Muller-Lyer illusion, their visual experience, *E*, must have impossible content, content that could be accurate in no metaphysically possible world seems to be something like this. On the one hand when we look at the Muller-Lyer illusion we definitely do successfully perceive the endpoints of each line, and the fact that the lines run between them. On the other hand, one line definitely appears shorter to us than the

other. Suppose we think, as those who believe that perceptual experiences carry content presumably will, that the explanation of all this is that *E* and other experiences like it both correctly represent the lines as each running between two the correct two points, and correctly represent the locations of those points, but also that they represent the lines as different in length. Then this content will not be accurate at any possible world, since there is no world where those lines run between just those points, and yet are of different length. Now, I have my doubts about the cogency of this reasoning. But I will set them aside here, and assume, for the sake of the argument, that Brewer is right that if experiences have representational content at all, then experiences of the Müller-Lyer illusion have impossible content. Instead, I want to focus on criticizing the reasons he gives for claiming that experiences cannot have impossible content.

Brewer's argument seems to be the following:

1) Conceivability is 'a guide to' possibility, when it comes to 'the macroscopic perceptible world' and this fact is not inexplicable.
2) One good explanation of this relies on the claim that perception provides us with the concepts we use to think about 'the macroscopic perceptible world.'
3) This explanation relies on perceptual experiences not having impossible contents.
4) There are no other good explanations of the phenomena mentioned in 1).

Conclusion: Perceptual experiences never have impossible content.

A first thought here might be that 3) is too strong since Brewer talks about what sort of perceptible experiences we can have as a 'guide' to possibility. But note that, if we don't have 3) but rather 3'): this relies on perceptual experiences very rarely having impossible content, then we cannot get the strong conclusion that no perceptual experiences ever have impossible content. Instead we can only get the weaker conclusion that the vast majority of perceptual experiences do not have impossible content. From this it clearly does not follow that in the

unusual case of perceptual experiences of the Muller-Lyer line, those experiences do not bear impossible content. So the argument needs 3) and not just 3').

Now, I think this argument is very weak, for several reasons.

Firstly, and most importantly, Brewer doesn't actually even sketch any explanation of the fact that conceivability is 'a guide to' possibility, in the (vaguely specified) domain of propositions[172] which are in some sense about medium sized ordinary objects and their properties. In fact, he doesn't even sketch an account of how exactly the concepts which we form about the perceptible properties we see are derived from perception[173], or cite anyone else's account of this from the literature[174]. Until we know exactly what the proposed account of concept acquisition is, let alone how this account is supposed to help explain the purported fact about conceivability being a guide to possibility it's hard to see why we should be persuaded by Brewer's argument.

Secondly, there may be trouble lurking for Brewer in the literature on whether there is a more general link between conceivability and possibility. On the one hand, some philosophers, such as Chalmers (2002), have drawn very general links between conceivability and possibility. Few of these philosophers have thought the link is anything like as simple as 'all conceivably true propositions are possibly true'. But they have generally taken the link to be completely general, and apply to all sorts of propositions, not just to propositions about the perceptible properties of medium-sized ordinary objects. Now, these philosophers have their own accounts of why

[172] I'll stick to doing things in terms of propositions rather than sentences here, though it's not clear to me which way of doing things Brewer prefers.

[173] There is some discussion of *what sorts* of encounters with perceptible properties allow us to form concepts which pick them out, in ch.5 of Brewer's (2011, pp.104-5), namely when we perceives objects which are paradigm exemplars of the properties in question. But it still contains no account of *how* this process works.

[174] Note that Brewer needs an account here that both does not entail that perception has representational content and yet is also such that it could be used to explain the link between conceivability and possibility in the domain of perceptible properties if experiences have representational content, providing only that that content is *never* inconsistent. If the account isn't consistent with any content view, then Brewer can rely on it alone without bringing in the Müller-Lyer line at all.

conceivability under certain circumstances will entail possibility[175]. The worry for Brewer's argument is that one of these accounts can be used to explain why conceivability is a guide to possibility in the domain of propositions he's interested in, rendering 4) above false. At the same time, many philosophers have doubted that there is any general link between conceivability and possibility at all. It is possible that the arguments they give against their being such a general link will also work against Brewer's claim that conceivability is a guide to possibility in the domain of propositions about medium-sized objects and their perceptible properties. Worse still for Brewer to the extent that we show that objections to their being a more general link between conceivability and possibility don't work, our confidence in their being a more general explanation of the link, which might cover the group of propositions Brewer is interested in will rise. So to some degree, attempts to solve the second problem here may exacerbate the first.

Thirdly and finally note that Brewer seems to need a very specific and strong link between conceivability and possibility within the domain he is interested in, and this looks particularly hard to justify. It is presumably not enough for Brewer to hold that, if a proposition about the perceptible properties of ordinary objects is conceivably true, then it is *very likely* metaphysically possible. It is hard to be sure of this, given that Brewer doesn't actually sketch the alleged explanation of the link between conceivability and possibility in the domain of perceptible properties. But presumably the explanation of the (purported) conceivability/possibility link which Brewer has in mind could only depend on their *never*, rather than rarely, being perceptual experiences with impossible content, if it was *never*, *r*ather than just rarely the case that conceivability failed to imply metaphysical possibility for propositions about perceptible properties. For suppose *most* of our perceptual experiences do not have impossible content. Then it is surely possible that most of our concepts which pick out perceptible properties of physical objects, are acquired in a manner consistent with Brewer's (shadowy) explanation of the alleged conceivability possibility link. But now, suppose that most of the propositions we

[175] See again Chalmers (2002).

entertain about medium-sized objects are ones we entertain by deploying the concepts which are acquired for Brewer in the right way, and *without deploying the concepts acquired from experiences with impossible content*. Presumably, we can use Brewer's explanation of the alleged conceivability possibility link to explain why for these propositions, conceivability implies possibility. But it is compatible with most of our perceptual experiences not having impossible content that a few, such as experiences of the Muller-Lyer illusion do. So in all likelihood Brewer doesn't just need the claim that, in general, if we can conceive of a particular scenario involving perceptible properties of ordinary objects, then it's probably possible. Rather he needs the stronger claim that conceivability always implies possibility here. Otherwise, Brewer's explanation of the conceivability possibility link will not rely on there being no perceptual experiences with impossible content, but merely on their being few, and 3) will be false.

I conclude then that Brewer's first argument is a failure. What about his second argument?

Brewer's second argument involves the worry that those who believe that experiences have content will, in the case of certain non-veridical perceptual experiences, end up committed to implausible views about which items those experiences make present to their subjects. Suppose that we think, as most who believe that perceptual experiences are content-bearers presumably will, that the representational content of an experience, either alone or along with facts about how that experience is caused, determines which objects the experience makes phenomenally present to its subject. Then, says Brewer, we can show that, given reasonable background assumptions, it will be possible for experiences to make items phenomenally present to me, even if the experiences grossly misrepresent them. The argument Brewer gives for this runs as follows. What objects an experience makes present to its subject is determined by either its singular or its existential content. If it is determined by its singular content, then the most obvious way for this to work is simply for experiences to make present to their subjects whatever objects enter into their singular content. But there seems to be no reason why an experience could not carry

singular content about an object, and yet still radically misrepresent that object. Meanwhile, in the case of the existential content of experiences, Brewer argues that the most natural way to do things is to pick out the object (or objects) that best fit the content in the immediate environment, perhaps with the added requirement that the object cause the experience 'in the right way'. Brewer argues that once again, there's no obvious reason why there's a limit to how badly an experience can misrepresent an object without making it present to its subject, on such a view. For best fit is relative not absolute, and being caused in the right way by an object doesn't impose any constraint on the range of existential contents an experience could posses. Brewer then goes on to argue that it is obviously not true that objects can be made phenomenally present to us by our experiences, regardless of how badly we misrepresent them, and that therefore the view that it is the representational content of experiences which determine which objects they make present to their subjects should be abandoned:

> Now, it seems to me that there are limits on the nature and extent of any errors involved in illusion, as opposed to hallucination, due to the fact that this involves genuine perceptual presentation of the object o that illusorily looks F. These are not hard and fast, there may be wildly abnormal circumstances in which the illusions I reject are in fact possible, and I certainly don't propose strict principles governing which are and which are not possible, or even suppose that there are any such general principles. Still, although a pair of lines of equal length about a metre away directly in front of me may look unequal in length, or may look unparallel when they are parallel, those very lines could not normally look like a single perfect circle, or two perfect circles, for example. Similarly, although a rabbit curled up on the chair next to me may look like a cat or a cushion, that very animal could not normally look like the Eiffel Tower. This is due to the fact that genuine perceptual presentation is in general incompatible with extreme error. (2011, p.73).

The first obvious thing to ask is whether the claim Brewer defends, that an experience cannot both misrepresent an object to a very high degree, and also make that object phenomenally present to its subject, is inconsistent with the SRAPC. The answer, sadly, is yes. In my account of illusion above, I committed to the claim that, provided we represent some property of an object accurately, that object will be phenomenally present to us[176]. But of course, representing some

[176] I was also neutral, rather than opposed to the idea that the objects an illusory experience misrepresents are always made phenomenally present to the subject of such an experience, even in the case where none

property of an object accurately, such as that it has a particular (uniform) colour across its whole surface, is compatible with radical misrepresentation of other of its properties, such as perhaps its shape[177]. Since the SRAPC is in conflict with the claim which Brewer takes to be intuitively obvious about the relationship between phenomenal presence and misrepresentation, I do need to give some kind of response to his argument here.

Now, at this point, I must say that I don't particularly share Brewer's intuition about the relationship between misrepresentation and phenomenal presence. It's not clear to me why it should be obviously impossible for an object to be phenomenally present to me, even though I badly misrepresent that object[178]. Still, merely reporting a difference of intuitions with Brewer here is not in itself especially interesting. So instead, I shall now argue instead that there is no obviously reliable source from which Brewer's intuition might hail, thus, I hope, casting sufficient doubt on it to discredit Brewer's second argument.

How could someone come to know that experiences which severely misrepresent their objects do not make those objects phenomenally present to their subjects? Presumably, this could not be known by introspection. For suppose that we could tell by introspection alone that a particular experience, which inaccurately represented the world to contain an object with particular properties did not make an object with those properties phenomenally present to one. One would thereby gain a way of telling certain experiences which were not successful perceptions, those which greatly misrepresent their objects, from successful perceptions with the same existential content. For the successful perceptions will make objects of the relevant type phenomenally present to their subjects. But of course, there is no reliable way to tell successful

---

of their properties are represented accurately. Though I think no trouble would result from switching from neutral to opposed on this stronger claim in the face of Brewer's argument.

[177] Or at least, there is nothing in the idea of representation which rules out that an experience both carry singular content about an object and also represent that object entirely inaccurately. This is clear possible in the case of thoughts with singular content, so there seems to be no reason to think that it is not also possible in the perceptual case. As Brewer himself points out, simply imposing some sort of limitation here for perception alone looks disturbingly ad hoc.

[178] Of course, the SRAPC does entail that if I misrepresent an object, at all, then it is not present to me bound together with (or present as instantiating, if we nominalize) that property.

perceptions from non-successful perceptions by introspection alone. Nor, clearly, has Brewer arrived at his claim that severe misrepresentation precludes making phenomenally present, on the basis of some kind of empirical investigation. And it's obviously not being defended on the grounds of parsimony or inference to the best explanation. Since it's not something which Brewer could come to know introspectively, and he hasn't come to know it empirically, and it's not being put forward on grounds of parsimony, how could he have come to know it?

Well, perhaps Brewer came to know it *a priori*, either wholly or in part. That is, perhaps it follows from some claims that are in some sense 'part of the concept' of phenomenal presence, that it's impossible for experiences to make objects which they misrepresent phenomenally present to their subjects. But this looks implausible, either on my account of 'phenomenal presence' or on Brewer's own. I introduced 'phenomenal presence' in ch.1, primarily as a term for the items associated with an experience which are the best candidates for being its qualia. I also added that talk of 'phenomenal presence' picks out a particular phenomenon which seems to be aptly describe by a particular metaphor of 'presence within experiences'. This phenomenon occurs in the case of perceptual experiences but not in the case of conscious thoughts. Nothing here seems to imply any conceptual connection between an experience's making an item phenomenally present to its subject, and its not being too severely misrepresented by that experience. Brewer's introduction of his notion of presence also does not exactly give it a rich role involving many connections to other notions:

> Something else that we take for granted is that physical objects are the very things that are presented to us in perception. It is extremely difficult to make this very natural idea precise. Indeed, a great deal of what follows effectively concerns various fundamental controversies concerning the notion of perceptual presentation. Still, I reserve the term throughout to express the utterly uncontested sense in which we see and otherwise consciously perceive physical objects: they are in this sense elements of perceptual consciousness. This claim that physical objects are presented to us in perception is intended as prior, and uncommitted, to any specific controversial theoretical elucidation of what such perceptual presentation consists in. (Brewer, 2011, p.2)

Admittedly Brewer might argue that the connection he makes here between presence and seeing is what secures the claim that an experience cannot both dramatically misrepresent an object and make that object phenomenally present to its subject. If so, however, I am happy to concede that in the cases where I have an illusory experience and radically misrepresent most properties of an object, whilst getting one or another exactly right, I do not 'see' that object. What Brewer has to demonstrate, in order to defend his claim via the connection between seeing things and their being phenomenally present to one, is that on any reasonable account of phenomenal presence, the claim that something can be phenomenally present to us only if we see it will hold. He does not attempt to do this.

I conclude then, that it's hard to see how Brewer's conviction that it's impossible for an experience to both make an object, *o* phenomenally present to its subject, and also severely misrepresent *o* can be justified. This suggests we should not give his claim that it is intuitively obvious that this is impossible too much weight. Not at the very least until he shows that the intuition is strong and widespread. Without this intuition, Brewer's second argument cannot even get started, so my defence of the SRAPC from the arguments of naïve realists who reject the view that experiences have representational content is now complete.

Section Four-A Difficulty for the SRAPC Regarding a Certain Possibility About Illusion

So far, I have argued that the SRAPC fairs better than other extant Phenomenal Difference views, and defended it from objections to various of its components which are found in the literature. It would be remiss of me however, not to mention that there is one *possible* difficulty with the SRAPC which I am not sure that I yet know how to solve. This difficulty involves the possibility that illusions are highly pervasive, so much so that we rarely represent the properties of objects around us accurately. I'll say why someone might think that this is the case, and then why if this is

the case, it causes problems for the SRAPC. I am not sure whether or not the problem for the SRAPC can be solved if illusion truly is very pervasive. I will suggest one way in which we might resolve the problem by modifying the SRAPC below. But it's not clear to me how tenable the modified theory is. The overall moral to be drawn from this chapter should, I think therefore be that the SRAPC is the best available view, conditional on both there being no way to solve the problems with the quasi-descriptivist view discussed in the previous chapter, and illusion turning out not to be extremely pervasive after all. If illusions are in fact extremely pervasive in the sense just discussed, *and* my attempt to modify the SRAPC below in order to account for this fails, then I think we have to drop the SRAPC altogether. Then it seems we should either look again at ways to resolve the problems with the quasi-descriptivist view discussed in ch.3 or adopt an Anti-Presence view. For the argument from pervasive illusion against the SRAPC clearly, if successful, works against other Phenomenal Difference views given the extended transparency thesis I argued for in ch.2 according to which successful perceptions make perceptible scenes phenomenally present to us.

So why might someone think that illusions are highly pervasive? Well, the main reason is that very often, how we perceive objects to be changes even though we have good reason to think that the objects themselves do not change, or at least do not change in their intrinsic properties. As Smith (2002, p.27) points out objects change their appearance with small changes in illumination, as for example, when objects gradually change how they appear to us as dusk falls. Likewise, objects change their appearance in respect to shape with changes in the angle of viewing (see Fish (2009, p.159) for some good examples). In all these cases, there is some pressure to say there is illusion involved. For there is some pressure to say a) that the colour or shape which the subject's experience represents the object to have has changed and b) that colours and shapes are intrinsic properties, and it is clear that the intrinsic properties of the objects need not have changed in the cases. To show that there is not illusion in such a case, we need to find some relational properties of the object, such that 1) they have changed with changes in illumination and viewing angle, and

2) plausibly, the change in how the object appears is a change in which of these properties it is represented to have. In the case of colours, this could lead to the denial of a) or b), since the view that colours are themselves relational properties is not absolutely outrageous; in the case of shape it seems more likely to lead to the denial of a). Fish tries to develop such an account in (2009, ch.6).

Furthermore, there are some reasons to think that such accounts involving successful perception of the relational properties of experiences are preferable to those which claim that such cases involve an illusory misrepresentation of the colours and shapes of objects. In the case of colour in particular, where the appearance of objects changes not just with illumination, but also with changes in the background against which the colour is viewed, it's hard to see what would make one set of viewing conditions those in which we saw the colour of an object as it is, and another not[179]. In addition there are more general reasons for thinking that pervasive error is unlikely. Suppose that most of the time, we misrepresented the shapes and colours of objects. Then there would be little connection between the typical cause of a mental state, or even its typical cause in ideal conditions[180], and what it represents. But it seems that, however difficult it is to work out an exact account of how to 'naturalize intentionality' there is probably some connection here[181].

Suppose however, that small changes in illumination or viewing angle do typically produce illusion. This would have dire consequences for the SRAPC. Since the SRAPC entails that in the case of illusions, the phenomenal presence facts are other than they seem, if, in most conditions of illumination and at many viewing angles, illusion occurs, then we get the consequence that things are almost never as they seem, with respect to what is phenomenally present to us. Furthermore, if changes in illumination, in particular, lead to illusion, then it seems that not only

[179] This point is also made by Fish (2009, p.156)).
[180] I'm assuming here that ideal condition can remain ideal even with minor changes in illumination and viewing angle.
[181] Though some, such as Millikan (1989) have tried to get an account of mental representation out of the notion of what properties it is the function of a mental state to indicate, rather than out of reliable tracking. It's not clear to me whether such views are likely to be more hospitable to the idea that shape perception involves pervasive misrepresentation, in the end.

will most of our experiences be partially illusory, but very many of our experiences will fail to represent the colour of any object accurately. The latter result will make the case where we are actually confronted with a perceptible scene, containing an appropriately tied together object and colour, a reasonably rare occurrence. And it will entail that, much of the time, none of the objects which seem to be present to us as bound together with certain colours, are in fact present to us as bound together with those colours. In effect, if colour and shape illusions are highly pervasive, then the proponent of the SRAPC will be forced to say that we are highly unreliable in judging what phenomenal properties our experiences have, even in ideal conditions. But at this point, presumably there will then be a lot of semantic pressure to interpret our phenomenal vocabulary and concepts in such a way that we are not so error prone, at least if there is any reasonable way of doing so. And it seems that if we adopt some sort of representationalist Anti-Presence view, then there will be a way of doing so. On such a view, the phenomenal properties of experiences supervene on their representational content, not the items they make present to their subjects. There is therefore no reason, on such views to think that we will be mistaken about the phenomenal properties of our experiences if we attempt to simply read them off the representational content of those experiences in the illusory or purely hallucinatory cases.

Furthermore, if illusions are highly pervasive in the sense just discussed, this argument will work equally well against any Phenomenal Difference view which also entails that introspection is misleading whenever we undergo illusory experiences. And this would seem to be a feature of any Phenomenal Difference view on which perceptible scenes are phenomenally present to us in the successful perceptual case. For illusions are not usually introspectively discernible from successful perceptions which bear the same existential content as they, without the aid of some background beliefs which would give us reason to regard a particular experience as illusory. So we know that, if illusions fail to make the same items phenomenally present to their subjects as would successful perception with the same existential content, this will not be noticeable on introspection. Hence a perceptible scene which is qualitatively the same will seem to be present

in the case where someone genuinely perceives an object *o* to be *P* and the case where someone merely has an illusory experience which misrepresents some object *o'* to be *P*. Furthermore, as we saw above there are good reasons for thinking that an illusory experience of an object as instantiating a property does not make present to its subject a perceptible scene which has as a part that object instantiating that property. So on any reasonable Phenomenal Difference view, we will indeed be disposed to make mistakes when we introspect illusory experiences. This means that if illusion is ubiquitous, it is Phenomenal Difference views in general, rather than the SRAPC specifically, that are in trouble, at least given extended transparency.

Note, by the way that, the matter discussed here is not an all or nothing one. The more pervasive illusion is, the more pressure this puts on Phenomenal Difference views, but it's not clear just how pervasive illusion has to be before those views are in serious trouble. For example, it's not clear how much trouble they are in if most of our experiences have some illusory element, but are largely accurate, in the sense that they mostly attribute properties to objects which they actually have. I am strongly inclined to think that if this is the case, Phenomenal Difference views, including the SRAPC specifically, can live to fight another day, in contrast with the case where most of the attribution of properties to objects by any randomly selected perceptual experience will attribute to those objects slightly different properties from those they in fact have. But it's hard to say exactly how to decide how much disposition to introspective error is too much. Note also that I have not ruled out that there may be other potential sources of extensive error apart from changes in illumination and viewing angle.

Now, there is one possible type of Phenomenal Difference view similar to the SRAPC which drops extended transparency, but which may be tenable even if illusions turn out to be extremely pervasive. In particular, recall that in Section Two above, I discussed how we should nominalize the SRAPC if there turn out to simply be no properties, and hence no perceptible scenes. And I suggested that nominalists should hold to a version of the SRAPC where for an experience to have

some determinate phenomenal property, *P*, is a matter of it making some object, *o* phenomenally present to its subject and *successfully* representing *o* to be a certain way, in contrast to the main formulation of the SRAPC on which having *P* is just a matter of making some perceptible scene or other item phenomenally present to a subject. Now because this nominalization of the SRAPC is phrased in terms of successful representation, rather than just representation, it will also be a view on which illusions will not have the phenomenal properties which they introspectively seem to have. For illusions will still differ from successful perceptions with the same existential content phenomenally, since these successful perceptions will involve only *successful* representation, and hence possess at least one phenomenal property the illusions fail to. But we can imagine a modified version of the view on which we drop the success requirement (indeed, I have already sketched such a view in footnote 21, back in Section Two). This leads to the view that for an experience, *E* to have some determinate phenomenal property *P*, is for it to make some object *o* present to its subject, and represent it to be *P*. This would still be a Phenomenal Difference view, since pure hallucinations make no objects phenomenally present to their subjects, and hence won't instantiate any phenomenal properties. But it's a view on which there is no reason at all to expect illusions to differ phenomenally from successful perceptions with the same existential content. Hence it's a view on which there is no reason to think that we will make mistakes about the phenomenal properties of illusory experiences when we introspect them. So it is not threatened if illusion turns out to be extremely pervasive.

I am reluctant to whole-heartedly endorse this view as the correct one to take if illusion turns out to be extremely pervasive, since I have not yet had time to explore the view fully. But at the very least, the view looks coherent at first glance. Hence there may be a Phenomenal Difference view which is quite similar to the SRAPC which fits perfectly well with most of our perceptual experiences being slightly illusory.

Bibliography

Adams, Robert M. 1979: 'Primitive Thisness and Primitive Identity,' *The Journal of Philosophy*, Vol.76, No.1, pp.5-26.

1981: 'Actualism and Thisness', *Synthese*, Vol.49, No.1, pp.3-41.

Alston, William 1999: 'Back to the Theory of Appearing', *Nous*, Vol.33, Supplement: Philosophical Perspectives 13, Epistemology, pp.181-203.

Bennett, Karen and McLaughlin, Brian 2011: 'Supervenience', in *Stanford Encyclopedia of Philosophy*, at http://plato.stanford.edu/entries/supervenience/

Block, Ned 1995: 'On a Confusion About a Function of Consciousness', *Behavioral and Brain Sciences*, Vol.18, No.2, pp.227-247.

1996: 'Mental Paint and Mental Latex', *Philosophical Issues*, Vol.7, pp.19-49.

2007: 'Consciousness, Accessibility, and the Mesh Between Psychology and Neuroscience', *Behavioral and Brain Sciences*, Vol.30, pp.481-548.

2010: 'Attention and Mental Paint', *Philosophical Issues*, Vol.20, pp.23-63.

2011: 'Perceptual Consciousness Overflows Cognitive Access', *Trends in Cognitive Science*, Vol.15, No.12, pp.567-575.

2014: 'Rich Conscious Perception Outside Focal Attention', *Trends in Cognitive Science*, Vol.18, No.9, pp.

and Fodor, Jerry 1972: 'What Psychological States Are Not', *Philosophical Review*, Vol.81, No.2, pp.159-181.

Brewer, Bill 2008: 'How to Account for Illusion', in *Disjunctivism: Perception, Action, Knowledge*, ed(s). Haddock and MacPherson, pp.168-180.

2011: *Perception and Its Objects*, Oxford University Press, Oxford.

Burge, Tyler 2005: 'Disjunctivism and Perceptual Psychology', in *Philosophical Topics*, Vol.33, No.1, pp.1-78.

Byrne, Alex 2001: 'Intentionalism Defended' in *Philosophical Review*, Vol.110, No.2, pp.199-240.

2006: ‘Colour and the Mind-Body Problem’, in *Dialectica*, Vol.60, No.60, No.3, pp.223-244.

2009: ‘Experience and Content’, in *Philosophical Quarterly*, Vol.59, pp.429-451

and Hilbert, David 2003: ‘Colour Realism and Colour Science’, in *Behavioral and Brian Sciences*, Vol.26, No.1, pp.3-21.

and Logue, Heather 2008a: ‘Introduction’ in *Disjunctivism: Contemporary Readings*, MIT Press, Cambridge Mass.

2008b: ‘Either/Or’, in *Disjunctivism: Perception, Action, Knowledge*, ed.Haddock, MacPherson, pp.57-92.

Byrne, Alex and Tye, Michael 2006: ‘Qualia Ain’t in the Head’, in *Nous*, Vol.40, No.2, pp.241-255.

Campbell, John 2002: *Reference and Consciousness*, Clarendon Press, Oxford.

Carruthers, Peter 1996: *Language, Thought and Consciousness: An Essay in Philosophical Psychology*, Cambridge University Press, Cambridge.

Chalmers, David 1996: *The Conscious Mind*, Oxford University Press, Oxford.

2002: ‘Does Conceivability Entail Possibility?’, in Hawthorne and Gendler (ed.s) *Conceivability and Possibility*, Oxford University Press, Oxford, pp.145-200.

2010: *The Character of Consciousness*, Oxford University Press, Oxford.

2011: ‘Verbal Disputes’, in *Philosophical Review*, Vol.120, No.4, pp.515-566.

Chisholm, Roderick 1942: ‘The Problem of the Speckled Hen’, in *Mind* Vol.51, pp.368-373.

Chudnoff, Elijah 2013: ‘Awareness of Abstract Objects’, in *Nous*, Vol.47, No.4, pp.706-726.

Cohen, Joshua 2002: ‘The Grand Grand Illusion Illusion’, *Journal of Consciousness Studies*, Vol.9, No.s5-6, pp.141-157.

Crane, Tim 2006: ‘Is There a Perceptual Relation?’ in *Perceptual Experience*, ed. John Hawthorne and Tamar Gendler, Oxford University Press, Oxford, pp.126-146

and Mellor, D.H. 1990: ‘There is No Question of Physicalism’, in *Mind*, Vol.99, No.394, pp.185-206.

Davidson, Donald 1970: ‘Mental Events’, in Foster and Swanson (ed.s) *Experience and Theory*, Duckworth, London.

1987: ‘Knowing One’s Own Mind’, in *Proceedings and Addresses of the American Philosophical Association*, Vol.60, pp.441-58.

Dennett, Daniel 1988: 'Quining Qualia' in *Consciousness in Modern Science*, ed.(s) A. Marcel and E. Bisiach, Oxfrod University Press, Oxford, pp.42-77.

1991: *Consciousness Explained*, Little, Boston.

Devries, Willem A. 1996: 'Experience and the Swamp Creatures', in *Philosophical Studies*, Vol.82, No.1, pp.55-80.

Dretske, Fred 1995: *Naturalizing the Mind*, MIT Press, Cambridge Mass..

1999: 'The Mind's Awareness of Itself', in *Philosophical Studies*, Vol.95, No.(s) 1-2, pp.103-124.

2004: 'Change Blindness', *Philosophical Studies*, Vol.120, pp.1-18.

Egan, Andy 2006: 'Appearance Properties', *Noûs*, Vol.40, No.3, pp.495-521.

Ellis, Jonathan 2010: 'Phenomenal Character, Phenomenal Concepts and Externalism', in *Philosophical Studies*, Vol.147, No.2, pp.273-299.

Fish, William 2009: *Perception, Illusion, Hallucination*, Oxford University Press, Oxford.

Fodor, Jerry 1990: *A Theory of Content and Other Essays*, MIT Press, Cambridge Mass.

Frey, Christopher 2013: 'Phenomenal Presence', in Uriah Kriegel (ed.), *Phenomenal Intentionality*, Oxford University Press, Oxford, pp.71-92.

Genone, James 2014: 'Appearance and Illusion', in *Mind*, Vol.123, pp.339-376.

Harman, Gilbert 1990: 'The Intrinsic Quality of Experience' in *Philosophical Perspectives*, Vol.4, pp.31-52.

Hawthorne, John 2004: 'Why Humeans Are Out of Their Minds', in *Noûs*, Vol.38, No.2, pp.351-358.

2009: 'Superficialism in Ontology', in Chalmers, Manley and Wasserman (ed.s), *Metametaphysics: New Essays on the Foundations of Ontology*, Oxford University Press, Oxford, pp.213-230.

Hill, Christopher S. 2009: *Consciousness*, Cambridge University Press, Cambridge.

Horgan, Terence, and Tienson, John 2002: 'The Intentionality of Phenomenology and the Phenomenology of Intentionality', in *Philosophy of Mind: Classical and Contemporary Readings*, Oxford University Press, Oxford, pp.520-533.

Horwich, Paul 1996: 'Comments on Dretske' in *Philosophical Issues*, Vol.7, pp.167-170.

Jackson, Frank 1975: 'On the Adverbial Analysis of Visual Experience', *Metaphilosophy*, Vol.6, No.2, pp.127-135.

1982: 'Epiphenomenal Qualia', *Philosophical Quarterly*, Vol.32, No.127, pp.127-136.

2004: ‘Representation and Experience’, in *Representation in Mind: new approaches to mental representation*, ed. H. Clapin, P. Slezak and P. Staines, Elsevier, Amsterdam, pp.107-124.

Johnston, Mark 2004: ‘The Obscure Object of Hallucination’, *Philosophical Studies*, Vol.120, No.1/3, pp.113-183.

Kennedy, Matthew 2009: ‘Heirs of Nothing: The Implications of Transparency’, in *Philosophy and Phenomenological Research*, Vol.79, No.3, pp.574-604

2013: ‘Explanation in the Good and Bad Cases’, in *Hallucination*, ed. Fiona MacPherson and Dimitris Platchias, MIT Press, Cambridge MA, pp.

Kind, Amy 2010: ‘Transparency and Representationalist Theories of Consciousness’, *Philosophy Compass*, Vol.5, No.10, pp.902-93.

King, Jeffrey C. 2007: *The Nature and Structure of Content*, Oxford University Press, Oxford.

Kriegel, Uriah 2007: ‘Intentional Inexistence and Phenomenal Intentionality’, in *Philosophical Perspectives*, No.21, pp.307-340.

2013: ‘The Phenomenal Intentionality Research Program’ in *Phenomenal Intentionality*, Oxford University Press, New York, pp.1-20.

Levine, Joseph 1983: ‘Materialism and Qualia: The Explanatory Gap’, in *Pacific Philosophical Quarterly*, Vol.64, pp.354-361.

Lewis, David 1973: *Counterfactuals*, Harvard University Press, Cambridge Mass.

1980: ‘Veridical Hallucination and Prosthetic Vision’, in *Australasian Journal of Philosophy*, Vol.58, No.3, pp.239-249.

1983a: ‘New Work for a Theory of Universals’, in *Australasian Journal of Philosophy*, Vol.61, No.4, pp.343-377.

1983b: ‘Individuation by Acquaintance and Individuation by Stipulation’, *Philosophical Review*, Vo.92, No.1, pp.3-32.

1986: ‘Against Structural Universals’, in *Australasian Journal of Philosophy*, Vol.64, No.1, pp.25-46.

Logue, Heather 2014: ‘Experiential Content and Naive Realism: A Reconciliation’, in Brogaard (ed.), *Does Perception Have Content?*, Oxford University Press, New York, pp.220-240.

Lycan, William G. 1996: *Consciousness and Experience*, MIT Press, Cambridge Mass.

Martin, M.G.F. 1997: ‘The Reality of Appearances’ in *Thought and Ontology*, ed. Mark Sainsbury, FrancoAngeli, Milan, pp.81-104.

1998: ‘Setting Things Before the Mind’, in *Current Issues in the Philosophy of Mind*, ed. O’Hear, A., Cambridge University Press, Cambridge, pp.157-80.

2001: 'Out of the Past: Episodic Recall as Retained Acquaintance', in *Time and Memory: Issues in Philosophy and Psychology*, ed. Christopher Hoerl and Teresa McCormack, Clarendon Press, Oxford, pp.257-284.

2002a: 'The Transparency of Experience', in *Mind and Language*, Vol.17, No.4, pp.376-425.

2002b: 'Particular Thoughts and Singular Thought', in *Logic, Thought and Language*, ed. O'Hear, A., Cambridge University Press, Cambridge, pp.173-214.

2004: 'The Limits of Self-Awareness', Philosophical Studies, Vol.120, No.s1-3, pp.37-89

2006: 'On Being Alienated', in *Perceptual Experience*, ed. John Hawthorne and Tamar Gendler, Oxford University Press, Oxford, pp.354-410.

Forthcoming: *Uncovering Appearances*

McDowell, John 1994: *Mind and World*, Harvard University Press, Cambridge Mass.

Milikan, Ruth 1989: 'Biosemantics', in *Journal of Philosophy*

Mole, Christopher 2011: *Attention is Cognitive Unison: An Essay in Philosophical Psychology*, Oxford University Press, Oxford.

Nagel, Thomas 1974: 'What Is it Like to Be a Bat?', *Philosophical Review*, Vol.83, No.4, pp.435-450.

Norman, Liam J., Heywood, Charles A. and Kentridge, Robert W. 2013: 'Object-Based Attention Without Awareness', *Psychological Science*, Vol.24, No.6, pp.836-843.

Nudds, Matthew 2009: 'Recent Work in Perception: Naive Realism and Its Opponents', in *Analysis*, Vol.69, No.2, pp.334-346.

Pautz, Adam 2007: 'Intentionalism and Perceptual Presence', in *Philosophical Perspectives*, Vol.21, pp.495-541.

2014: 'The Real Trouble With Armchair Arguments Against Phenomenal Externalism', in Sprevak and Kallestrup (ed.s) *New Waves in Philosophy of Mind*, Palgrave-MacMillan, Basingstoke, pp.153-181

Peacocke, Christopher 1983: *Sense and Content*, Clarendon Press, Oxford.

1992: *A Study of Concepts*, MIT Press, Cambridge, Mass..

Price, Huw 2011: *Naturalism Without Mirrors*, Oxford University Press, Oxford.

Prinz, Jesse 2012: *The Conscious Brain*, Oxford University Press, Oxford.

Pritchard, Duncan 2012: *Epistemological Disjunctivism*, Oxford University Press, Oxford.

Rensink, Ronald A., O'Regan, J. Kevin, and Clark, James J. 1997: 'To See or not to See: The Need for Attention to Perceive Changes in Scenes', *Psychological Science*, Vol.8, pp.368-373.

Schellenberg, Susanna 2011a: 'Perceptual Content Defended', *Noûs* Vol.45, No.4, pp.714-750.

2011b: 'Ontological Minimalism About Phenomenology', in *Philosophy and Phenomenological Research*, Vol.83, No.1, pp.1-40.

Schroer, Robert 2012: 'Representationalism and the Scene-Immediacy of Visual Experience: A Journey to the Fringe and Back', in *Philosophical Psychology*, Vol.25, No.4, pp.595-615.

Schwitzgebel, Eric 2006: 'Do Things Look Flat', in *Philosophy and Phenomenological Research*, Vol.72, No.3, pp.589-599.

Searle, John 1958: 'Proper Names', in *Mind*, Vol.67, pp.166-173.

1983: *Intentionality*, Cambridge University Press, Cambridge.

Shoemaker, Sydney 1980: 'Causality and Properties', in (ed.) Peter van Inwagen, *Time and Cause*, D. Reidel Publishing Company, Dordrecht, pp.109-135.

2000: 'Introspection and Phenomenal Character', in *Philosophical Topics*, Vol.28, No.2, pp.247-273.

2006: 'On the Way Things Appear', in John Hawthorne and Tamar Gendler (ed.) *Perceptual Experience*, Oxford University Press, Oxford, pp.461-480.

Sider, Theodore, 2011: *Writing the Book of the World*, Oxford University Press, Oxford.

Siegel, Susanna 2004: 'Indiscriminability and the Phenomenal', in *Philosophical Studies*, Vol.120, No.1, pp.91-112.

2010: *The Contents of Visual Experience*, Oxford University Press, Oxford.

Simons, D.J. 2000: 'Current Approaches to Change Blindness', *Visual Cognition*, Vol.7, No.s1-3, pp.1-16.

Smith, A.D. 2002: *The Problem of Perception*, Harvard University Press, Cambridge Mass.

Soames, Scott 1987: 'Direct Reference, Propositional Attitudes and Semantic Content', *Philosophical Topics*, Vol.15, No.1, pp.47-87.

2010: *What is Meaning*? Princeton University Press, Princeton.

Stich, Stephen 1999: *Deconstructing the Mind*, Oxford University Press, Oxford.

Stoljar, Daniel 2004: 'The Argument from Diaphanousness', in Escurdia, Stainton, and Viger (ed.s) *Language, Mind and World: Special Issue of the Canadian Journal of Philosophy*, Alberta University Press,

Sturgeon, Scott 2000: *Matters of Mind*, Routledge, New York.

Tegmark, Max 2014: *Our Mathematical Universe: My Quest for the Ultimate Nature of Reality*, Allen Lane, London.

Travis, Charles 2004: 'The Silence of the Senses', *Mind*, Vol.113, pp.57-94

2013: *Perception: Essays After Frege*, Oxford University Press, Oxford.

Tye, Michael 1995: *Ten Problems of Consciousness*, MIT Press, Cambridge Mass.

2000: *Consciousness, Color and Content*, MIT Press, Cambridge Mass.

2002: 'Representationalism and the Transparency of Experience', *Noûs*, Vol.36, No.1, pp.137-151.

2009: *Consciousness Revisited: Materialism Without Phenomenal Concepts*, MIT Press, Cambridge Mass.

2014: 'Transparency, Qualia Realism and Representationalism', *Philosophical Studies*, Vol.170, No.1, pp.39-57.

Forthcoming: 'What is the Content of a Hallucinatory Experience?', in B. Brogaard (ed.) *Does Perception Have Content?* Oxford University Press.

Williamson, Timothy 2007: *The Philosophy of Philosophy*, Blackwell Publishing, Oxford.