

Perceiving What Must Be; A New Argument in Favour of Naïve Realism.

Mr Nicholas Thomas Courtney

PhD

University of York

Philosophy

January, 2020

Abstract

I offer an original argument in favour of Naïve Realism, over Intentionalism. I shall argue that we visually experience: 1. Spatial composition relations, and 2., the metaphysical necessity of those composition relations. Naïve Realism is able to accommodate both 1. and 2., in a way that is consistent with our visual phenomenology. Representationalist versions of Intentionalism are unable to accommodate even 1., in a way that is consistent with our visual phenomenology, and other forms of Intentionalism have some trouble accounting for 1., in a way that is consistent with our visual phenomenology. No form of Intentionalism is able to accommodate 2., in a way that is consistent with our visual phenomenology. I will also develop the Naïve Realist account such that it is, for the first time, able to provide a satisfactory explanation of our perceptual experience of objects (as well as of properties).

Author's Declaration

I declare that this thesis is a presentation of original work and I am the sole author. This work has not previously been presented for an award at this, or any other, University. All sources are acknowledged as References.

Table of Contents

<i>Introduction – Moving Beyond Transparency</i>	7
<i>Chapter 1 – Groundwork; The ‘Phenomenal Contents’ of Perceptual Experience.</i>	26
1.1 Three Theory Neutral Theses Regarding Perceptual Content	26
1.2 Summary	37
<i>Chapter 2 – Naïve Realism and Perceptual Experience of Objects.</i>	38
2.1: Why a Perceptual Relation with <i>Objects</i> ?.....	38
2.2 - The Presentation of Objects <i>as Such</i>	51
<i>Chapter 3 – Implications and Empirical Support.</i>	59
3.1: The Empirical Basis.....	59
3.2 - Perceptual Experience of only one Object at a Time.	63
3.3: Cases in which we do not Select an (Single) Object.	65
<i>Chapter 4 - Perceptual Experience of Spatial Composition</i>	69
4.1. The Case for Spatial Composition in Phenomenal Content (Claim 1)	71
4.1.1 Argument 1 - Perceptual Experience of Similarity Relations.....	74
4.1.2 Argument 2 – Perceptual Experience of the <i>Ways</i> in Which Shapes are Similar	78
4.1.3 Argument 3 – Our Capacity to Purposefully Manipulate Shapes.....	82
4.2 An Alternative Explanation	85
4.2.1 An Objection – Merely Perceiving- <i>That</i> the Square is Composed of the Straight Lines.....	86
4.2.2 Perceptual Uses and Evidential Uses of the Language of Appearing.....	90
4.3. The Necessity of Phenomenal-Composition (Claim Two), and Ruling out Cognitive Penetration (Claim 3).....	96
4.4. Kanizsa Shapes: A Counter-Example to Claim Two?	102
4.5. Summary.....	111
<i>Chapter 5 – Knowledge of Spatial Composition.</i>	113
5.1 Analytic Truths.....	114
5.1.1 The Relevance of Analyticity.....	114
5.1.2. Williamson, and General Problems with the Epistemological Notion of ‘Analytic Truths’.	116

5.2. Spatial Composition Relations as Analytic Truths	129
5.2.1. Assent and Justification.....	129
5.2.2 How Might the Understanding-Assent Links Hold? Constitutive Inter-Conceptual Relations.	130
5.2.3. Undermining Inter-Conceptual Relations Among Shape Concepts.....	134
5.3. Innate Beliefs	137
5.4 Summary.....	141
 <i>Chapter 6 - Knowledge of the Metaphysical Necessity of Spatial Composition.143</i>	
6.1. The Source of our Knowledge of the <i>Metaphysical Necessity</i> of Spatial Composition Relations.....	144
6.1.1 Induction.....	145
6.1.2 Innate Beliefs, Take Two.....	147
6.1.3. Conceivability.....	148
6.1.4 Knowledge of Counter-Factuals.	154
6.2. A Positive Proposal	160
6.2.1 Knowledge of Modality Originates in Knowledge of Essence	160
6.2.2. – The Epistemic Role of Acquaintance	165
6.3. Objection – Revelation.....	168
6.4. Summary.....	173
 <i>Chapter 7 – Causal Accounts of Content Fixation. 175</i>	
7.1. Causal Accounts of Content Fixation	176
7.2. The Composition Objection	187
7.3. Responses.....	192
7.3.1 An Additional Element of Phenomenal Character	192
7.3.2 The Analyticity of Spatial Relations.....	194
7.3.3 Causal Over-Determination.....	197
7.4. The Modality Objection	202
7.5. Conclusion.....	205
 <i>Chapter 8 – Consumer Semantics and Benefit-Based Accounts of Content Fixation. 206</i>	
8.1 – Consumer Semantics and Spatial Composition: The Basics.....	207
8.2 –Perceptual Content for Believers	210

8.3 – A Commitment to Conceptualism.....	216
8.4 – Price’s Benefit-Based Account, The Composition Objection, and The Modality Objection.....	219
8.5 Conclusion.....	225
<i>Chapter 9 - Naïve Realist Accounts of Content Fixation.....</i>	<i>227</i>
9.1. The Naïve Realist Account of Perceptual Experience of Spatial Composition, and of its Metaphysical Necessity.....	227
9.2. A Further Advantage of Naïve Realism: Harmony of Content and Phenomenal Character.	232
<i>Thesis Conclusion</i>	<i>239</i>
<i>References.....</i>	<i>242</i>

Introduction – Moving Beyond Transparency

I intend to argue in favour of a Naïve Realist account of *visual* experience, over Intentionalist accounts of visual experience¹. I will begin with a characterisation of the relevant views. These views are different views about what visual *experience* – a type of event – *fundamentally* is. As Martin recognises, ‘we need some characterizations of events which reflect their nature or what is most fundamentally true of them’ (2006: 360). This is what a philosophical account of perceptual experience seeks to describe. Accounts of perceptual experience differ because they offer different characterizations of what is most fundamentally true of perceptual experiences. But a little more needs to be said about exactly what this amounts to. Following Martin:

‘I will just assume for the sake of this discussion that we can make sense of the idea that there are some privileged classifications of individuals, both concrete objects and events, and that our talk of what is essential to a given individual tracks our understanding of the kinds of thing it is. That is, I will assume the following: entities (both objects and events) can be classified by species and genus; for all such entities there is a most specific answer to the question, ‘What is it?’ In relation to the mental, and to perception in particular, I will assume that for mental episodes or states there is a unique answer to this question which gives its most specific kind; it tells us what essentially the event or episode is. In being a member of this kind, it will thereby be a member of other, more generic, kinds as well. It is not to be assumed that for any description true of a mental event, there is a corresponding kind under which the event falls’ (2006, 361).

With this elucidation of what an account of perceptual experience seeks to describe in

¹ I will often use ‘perception’, and ‘perceptual experience’, to refer to vision or visual experience. If I wish to refer to perception, or perceptual experience, *in general*, then I will be explicit about this.

hand, we may now distinguish the relevant views as follows²:

Naïve Realism/Relationalism: perceptual experience is, *fundamentally*, either a primitive relation between a subject and the mind-independent objects, properties and relations (directly) perceived, or (in cases of hallucination/illusion) an experience indistinguishable from this.

Intentionalism: Perceptual experience is, *fundamentally*, a mental state that is *about* mind-independent objects, properties and relations. When the subject is not subject to hallucination or illusion the subject thereby (directly) perceives those mind-independent objects, properties and relations.

A further crucial difference between these views will become apparent when, in chapter 1, we consider what Naïve Realists and Intentionalists might say about the fixation of perceptual content. In anticipation, Naïve Realism is distinguished from Intentionalism by the special role that the subject's immediate environment plays, for the Naïve Realist, in the fixation of the subject's perceptual content.

The last 30 years of philosophy of perception have been dominated by intense discussion of the putative transparency of experience. It is widely accepted as an accurate phenomenological observation, and many prominent philosophers of perception have believed that substantive metaphysical theses can be motivated by the transparency of experience. But once one delves into the literature it becomes clear that there is almost no agreement about what the implications of the transparency of experience really are. To this day, the debate rages on, without any sign of a convergence of ideas. Indeed, the only thing that all authors writing on the topic seem to agree on is that the transparency of experience is something worth talking about.

I believe that this entire discussion has been on the wrong track, and that a fixation on the transparency of experience (TE) has stymied progress in the philosophy of

² These definitions are exclusive, but not exhaustive, for there are other views of perceptual experience available (Sense Data Theory, for example).

perception. In this introduction I will give the dialectical situation surrounding TE some consideration, and I will offer a diagnosis of where and how it has gone wrong. It will be found that there are two principle problems with arguments that utilise TE, and these two problems explain why the TE is powerless to cast light on the kinds of questions that philosophers have hoped that it would cast light on. Then, after the groundwork has been laid in chapter 1, and once I have developed a positive Naïve Realist proposal in chapters 2 and 3, in chapter 4 I will offer some original phenomenological observations. These new phenomenological observations will be instrumental in my argument in favour of Naïve Realism, and over Intentionalism, which I present in chapters 7-9.

Allen (2016) offers the following recent construal of the transparency of experience (TE):

'The claim that visual experience is transparent can be understood as the conjunction of a positive and a negative thesis. The positive thesis is that reflection on our visual experiences reveals that we are aware of mind-independent objects, their properties, and relations. The negative thesis is that reflection on our visual experiences reveals awareness of nothing else—in particular, we are not aware of any entities distinct from the mind-independent objects and their properties that populate our environment' (2016, 13)

Various Naïve Realist have endorsed TE: Martin (2002), Campbell (2005) and Allen (2016). But Intentionalists have also appealed to TE. For example, Representationalists, Harman (1990) and Tye (1992, 1995), and Phenomenal Intentionalists, Horgan & Tienson (2002) and Horgan et al. (2004)³, all appeal to TE. The fact that the advocates of numerous competing views have all appealed to TE in support of their favoured view ought to be cause for doubt over the utility of TE.

³ Horgan & Tienson (2002), and Horgan et al. (2004) are not explicit in their appeal to TE. Nonetheless, Bordini (2016) interprets them as doing so.

Moreover, although TE is widely endorsed, it is by no means universally accepted. Block (1996) argues that TE is a false phenomenological claim. Block (1996) points to pain and orgasms as examples of features of experience that do not even *seem* to be features of the mind-independent environment⁴. Of course, our concern is with *visual* experience, and one might hold that only visual experience is transparent in the way described above. In which case Block's examples are not relevant. In response, the opponent of TE might postulate phosphene-experiences (Block: 1996), or the blurriness of blurry vision (Kind: 2003), as counter-examples to TE⁵. Counter-examples, such as these, constitute the first 'principle problem' with arguments that utilise TE, which I alluded to earlier.

Given what has been said, it is tempting to jettison TE at the outset, but I think that some further consideration of the dialectical situation surrounding TE will be of value to us. For once we have a greater understanding of where TE arguments have gone wrong, I will be in a position to explain how my own phenomenologically based argument will transcend the limitations of TE. As my aim in this thesis is to argue in favour of Naïve Realism, I will consider the use of TE in arguments in favour of this position. But TE was first offered in support of Representationalism – a specific variant of Intentionalism -, and since then it has continued to be most closely associated with Representationalism, so when I talk, in this Introduction, about the ways in which TE has been used to motivate Intentionalism, I will focus on the use of TE to argue in favour of Representationalism.

Recall, Intentionalists believe that a perceptual experience is, fundamentally, a mental state that is *about* the objects, properties and relations perceived.

⁴ Block (1996) refers to features like orgasms and pain as 'mental latex'. These features are to be distinguished from 'mental paint'. Mental paints, like the pigments of real paint, serve to represent (potentially mind-independent) properties, and therefore might introspectively seem to *be* mind-independent properties (even though they are not).

⁵ In Block's (1996) view phosphene-experiences involve mental paint that is usually used to represent colour, but which is not being used to represent colour, and which does not even introspectively *seem* to be being used to represent colour, or to be colour. Conversely, Block would presumably hold that the blurriness of blurry vision pertains to mental latex, because blurriness can qualify visual experiences of anything and everything.

Representationalism is the view that the ‘phenomenal character’ of perceptual experience – what it is like to undergo a perceptual experience – is determined exhaustively by what the perceptual experience is *about*. This view is often contrasted with the view that perceptual experience involves ‘qualia’, where ‘qualia’ are intrinsic (and usually non-physical) properties of an experience, which determine what it is like to undergo the experience. Below is a prominent example of an argument from TE, against qualia, and in favour of Representationalism:

‘Standing on the beach in Santa Barbara a couple of summers ago on a bright, sunny day, I found myself transfixed by the intense blue of the Pacific Ocean. Was I not here delighting in the phenomenal aspects of my visual experience? And if I was, doesn’t this show that there are visual qualia?’

I am not convinced . . . I experienced blue as a property of the ocean not as a property of my experience. My experience itself certainly wasn’t blue. Rather it was an experience that represented the ocean as blue. What I was really delighting in, then, were specific aspects of the content of my experience. It was the content, not anything else, that was immediately accessible to my consciousness and that had aspects that were so pleasing’ (Tye, 1992, 160).

The above argument is unsatisfactory. And Tye’s presentation actually makes the problem with the argument very clear. For an Intentionalist, the ‘content’ of a perceptual experience is the accuracy conditions that are associated with that perceptual experience. But ‘perceptual content’, as the Intentionalist understands this, is something possessed by perceptual states and not by the objects of perception. So Tye’s suggestion that ‘content’ is immediately accessible to his consciousness actually violates TE.

Even when Representationalists do not explicitly claim that content is immediately accessible to consciousness, it seems that they will still be committed to this. The Representationalist believes that phenomenal character is identical with a certain kind of content (or that it supervenes on a certain kind of content). It is usually assumed that we are (or at least can be) introspectively aware of phenomenal

character. But, once again, according to TE we only seem to be introspectively aware of the objects and properties that we perceive. I suggest that the reason that transparency arguments for Representationalism can seem persuasive is because it is easy to elide perceptual contents with the objects and properties that are ostensibly represented in perception. But when we recognise that these are distinct then we see that these arguments are unsuccessful.

Naïve Realism, as I understand it, holds that (at least part of) the phenomenal character of perceptual experience is identical with the objects and properties perceived. There is a very simple argument that will take us from TE to Naïve Realism. On the assumption that we are aware of the phenomenal character of our perceptual experiences (or that we at least can be), the claim that in perceptual experience we are not aware of anything other than 'the mind-independent objects, their properties, and relations' entails that phenomenal character is identical to those mind-independent objects, their properties, and relations. This implies that if we take experience at face value – i.e. if we assume that it is not misleading – then Naïve Realism must be true. More recently, Tye (2009, 2015) and Speaks (2015) seem to have become aware that Representationalism is vulnerable to arguments of this form, and they have offered responses.

Tye (2009, 2015) now claims that phenomenal character is not after all identical with content, as he used to say (or, more specifically, to PANIC content (Tye: 1995)), but that it is instead identical with the property-types that are represented in that content:

‘Consider again what the thesis of transparency tells us. It tells us that in the case of perceptual experiences, the only qualities of which we are introspectively aware are qualities of external things if they are qualities of anything at all. But intuitively, we are aware of phenomenal character when we introspect. The conclusion to draw is that the phenomenal character of a perceptual experience consists in, and is no more than, the complex of qualities the experience represents. Thus, the phenomenal character of the experience of red just is red. In being aware of red, I am aware of what it is like

to experience red, since what it is like to experience red is simply red... phenomenal character is not the same as representational content' (2009: 119).

'The only properties of which I am aware are external properties, including the color red. So, the phenomenal character of my experience is a cluster of external properties. In the simplest case, it is just the color red' (2015: 484).

'What matters to the phenomenology (phenomenal character) is the cluster of properties represented by the experience (see Tye 2014a and 2014c). This cluster of properties is not itself a content at all' (2015: 486).

Comparing this view to his previous one, he says 'my earlier work was not transparent enough about transparency' (2015: 485).

The difference between Tye's current view and Naïve Realism is that Tye is not claiming that the particular perceived object and its properties are identical with any elements of phenomenal character. This is made quite vivid by his account of hallucination:

'In (non-veridical) hallucination, on my view, I am confronted with a property that is not locally instantiated even though I experience it as such. Some philosophers profess to be very puzzled as to how this is possible. I profess to be puzzled at their puzzlement. Suppose I hallucinate something red. In doing so, I am aware of the color red. This is a matter of my undergoing an experience that represents the color red—an experience, that is, of a type that, under Normal conditions, tracks red (in first approximation). There is, then, a complex relation obtaining between my token experience as I hallucinate (given that there are token experiences) and the color red. This is what grounds my de re awareness of the property even though there is no local instance of it. What is so puzzling about that?' (2015, 485, footnote 3).

Tye is vulnerable to the objection that, introspectively, we do not seem to be merely aware of property types, but also of particular objects instantiating particular property instances (or tropes). Indeed, Tye's own description of his perceptual

experience, standing on the beach in Santa Barbara, has this implication. For he says ‘I experienced blue as a property of the ocean’. And the TE descriptions offered by other Representationalists invoke more paradigm examples of objects. For example, Harman says ‘When Eloise sees a tree before her, the colors she experiences are all experienced as features of the tree and its surroundings’ (1990, 39). Yet this is something that Tye seems to be powerless to explain. In chapters 2 and 3 we shall consider how the Naïve Realist ought to account for our perceptual experience of objects (and of particularity in general)⁶.

Speaks (2015) responds to the above argument, in favour of Naïve Realism, in a different way. He maintains that phenomenal character is a property of perceptual experience, as Intentionalists ordinarily suppose it to be, but he claims that phenomenal character is not introspectively accessible. So, with Tye, Speaks is able to maintain that what we are introspectively aware of in perceptual experience is limited to what is represented by the perceptual experience. But on Speaks’ view phenomenal character becomes an entirely theoretical notion, for if we cannot attend to it then we cannot refer to it ostensively. Yet Speaks offers no theoretical role for phenomenal character to play. So it seems that the notion has been emptied of all significance. Indeed, towards the end of his (2015) Speaks expresses some doubts over the utility of the notion. Contra Speaks, I think that Shoemaker (another author who makes frequent appeals to the Transparency of Experience) is right when he says:

‘If anything deserves to be called the phenomenal character of our experiences, it is the part of their *introspectable* nature that reflects how things look, feel, taste, smell, or sound to us’ (italics added) (2000, 258).

I will assume that if we are to talk of ‘phenomenal character’ at all then it had better be something introspectively accessible. In chapter 1 I will introduce a notion of ‘phenomenal character’ on which phenomenal character is introspectively accessible.

⁶ Unfortunately, I have not the space in this thesis to consider how the Intentionalist might account for our perceptual experience of objects (and of particularity in general). I believe that Intentionalists of all varieties face a serious problem accounting for this in a way that is consistent with our phenomenology. Previous drafts of this thesis included a chapter dedicated to making this argument.

So Speaks is wrong. If TE is a correct phenomenological observation, and if introspection is not misleading in that regard, then it seems that the Intentionalist must make the move that Tye makes⁷.

However, I think that TE is too strong. Let's return to Allen's construal of TE:

'The claim that visual experience is transparent can be understood as the conjunction of a positive and a negative thesis. The positive thesis is that reflection on our visual experiences reveals that we are aware of mind-independent objects, their properties, and relations. The negative thesis is that reflection on our visual experiences reveals awareness of nothing else—in particular, we are not aware of any entities distinct from the mind-independent objects and their properties that populate our environment' (2016: 13)

The negative component of TE has come under pressure from counter-examples of the kind offered by Block (1996) and Kind (2003). But the positive component of TE may also be too strong, and this is the second principle problem with arguments that utilise TE. The positive component can be interpreted as implying that we are aware of objects and their properties *as* mind-independent; that is, as implying that the *mind-independence* of objects and their properties shows up in the phenomenology of perceptual experience. But it is far from clear that perceptual experience is anything but silent on the metaphysical status of the objects and properties that we perceptually experience. That is, it is not obvious that the objects and properties that we perceptually experience wear their mind-independence on their sleeve, so to speak. More formally, it is not obvious that what presents itself to introspection is *manifestly* mind-independent.

There is another reading of the positive component of TE, on which TE does not have the above (contentious) implication. Allen says that '*reflection* on our visual experiences reveals that we are aware of mind-independent objects, their properties,

⁷ Actually, if we stick with Allen's formulation of TE then that won't be enough. For Allen's formulation entails that *objects* show up in perceptual experience, and, as I have already mentioned, it seems that Tye faces problems in this regard.

and relations' (italics added). Fastening on Allen's use of the word 'reflection', the positive component of TE could be interpreted as merely implying that, after reflection on our experience, we would be inclined to *judge* that (in visual experience) we are aware of *mind-independent* objects, properties and relations. In which case, this is not a purely phenomenological datum.

In fact, if the above interpretation is correct then the positive component of TE collapses into the claim that it is the *common sense* view that, in visual experience, we are aware of *mind-independent* objects, their properties, and relations. But this common sense view could simply be due to some cognitive bias. And while it is always possible that our introspective judgements are influenced by such biases, in so far as TE is meant to count as evidence that we really *are* aware of mind-independent objects and properties, the assumption must surely be that we are not, in this case, misled by any such bias. So in order to keep the various factors straight, it seems that we should interpret TE as a *purely* phenomenological datum, albeit a defeasible one (*defeasible* because it remains possible that we are in fact subject to some kind of cognitive bias).

Campbell (2014) sets forth a detailed and complex explanation of how, consistent with our phenomenology, objects might present themselves in perceptual experience *as* mind-independent. But the explanation that Campbell offers requires that we *first* endorse Naïve Realism. In which case, TE cannot be used as a theory neutral phenomenological observation with which we might motivate one view of perceptual experience (Naïve Realism, for example) over the others. So given that we are interested in the question of whether Naïve Realism or Intentionalism is the correct view of perceptual experience, we can hardly appeal to Campbell's (2014) argument to support the transparency of experience.

Martin (2002) seems to concede that the positive component of TE is under-motivated. Talking about how the Sense Data theorist might respond to the Tye passage quoted above, Martin admits that they are free to say:

'How is he so sure that it is the Pacific Ocean that he delights in when he turns his attention inwards, and not some mind-dependent blue expanse similar in

character to how Tye takes the Pacific to be? After all, the response might go, how could introspection alone show that the objects and entities that Tye can identify must be mind-independent, physical objects. The objector may concede that we typically are inclined to believe that we are presented with mind-independent objects in experience, but what they question is whether that belief can be adequately supported by introspection of experience alone' (2002: 382-382).

But he sets the objection to one side, saying of TE 'I am less concerned in this paper with how the sense datum theorist can respond to the challenge than how a defender of an intentional view should develop it' (2002: 382-382). I, too, am more concerned with the dispute between the Naïve Realist and the Intentionalist, so let's turn to what Martin says in this regard.

As Soteriou (2016) puts it, the core of Martin's (2002) argument against Intentionalism is that certain phenomenal features of experience constitute the apparent presentation of an object, and these features introspectively *seem* to be sufficient for the existence of just such an object. The Naïve Realist says that the phenomenal features really are sufficient for the existence of just such an object. And of hallucinations the Naïve Realist says that it merely *seems* as though one's experience involves those phenomenal features that really are sufficient for the existence of the object.

In contrast, the Intentionalist is committed to a view on which those features of experience that constitute the apparent presentation of an object are never sufficient for the existence of any such object. This is why the Intentionalist is able both to maintain that presented objects are mind-independent, and also to maintain, against the Naïve Realist, that one may be in a mental state of the same kind – involving exactly the same phenomenal features - whether or not there is a mind-independent object there to be perceived.

Martin's argument is that if, as he claims, those phenomenal features of experience that constitute the apparent presentation of an object really do introspectively *seem* to be sufficient for the existence of just such an object then introspection supports

views on which those phenomenal features really are sufficient for the existence of such an object. In which case, introspection supports Naïve Realism over Intentionalism. However, this argument does not actually require anything so strong as TE, because Martin does not claim that the object introspectively *seems* to be mind-independent, and the argument requires no such assumption (it would only require something so strong as *that* if it were being used as an argument against Sense Data Theory). The problem for the Intentionalist owes only to the fact that the phenomenal features introspectively seem to be sufficient for the existence of the object of perceptual experience, be this object *mind-independent* or *mind-dependent*.

If there is a transparency claim doing any work in Martin's argument then it seems to be something of the following form:

Weak Transparency (WT): when we introspect our visual experience we find but a single set of properties. We may think of these properties either as mind-dependent or as mind-independent. What we do not find in visual experience is duplicate sets of properties, regarding which we might think of the one set as mind-dependent (and presenting), while we think of the other set as mind-independent (and presented).

In other words, when we introspect visual experience we find only the objects (in the most general sense of the word 'object') of perceptual experience. WT speaks against views on which perceptual experience is only sufficient for some set of (mind-dependent) properties that serve to *present* the objects of perception, and on which perceptual experience is *not* sufficient for those objects of perceptual experience. The implication of WT is that a perceptual experience really is *sufficient* for the objects of the perceptual experience (be these objects mind-dependent or mind-independent).

WT has the advantage of avoiding the second principle problem with arguments that utilize TE, for it does not imply that the objects and properties of perceptual experience are *manifestly* mind-independent. Shortly I will come back to WT, and to how something like it might be used in a variant of Martin's argument in favour of Naïve Realism. First I wish to consider two further transparency claims, identified by

Nida-Rumelin (2007), and to consider what the implications of these other claims are, and how these claims relate to WT. Nida-Rumelin is interested to demonstrate that no plausible transparency claim supports Representationalism over other versions of Intentionalism which (contra-Representationalism) do acknowledge that experiences have *intrinsic* properties that we can introspectively attend to (she calls such properties '*intrinsic* phenomenal character'). Nida-Rumelin quotes the following passage from Tye:

'Suppose you are facing a white wall, on which you see a bright red, round patch of paint. Suppose you are attending closely to the color and shape of the patch as well as the background. Now turn your attention from what you see out there in the world before you to your visual experience. Focus upon your awareness of the patch as opposed to the patch of which you are aware. Do you find yourself suddenly acquainted with new qualities, qualities that are intrinsic to your visual experience, in the way that redness and roundness are qualities intrinsic to the patch of paint? According to some philosophers, the answer to this question is a resounding "No". As you look at the patch you are aware of certain features out there in the world. When you turn your attention inwards to your experience of those features, you are aware that you are having an experience of a certain sort but you are aware of the very same features; no new features of your experience are revealed. In this way your experience is transparent or diaphanous. When you try to examine it you see right through it, as it were, to the qualities you were experiencing all along in being a subject of the experience, qualities your experience is of'. (Tye, 2003: section 6)

Nida-Rumelin claims that Tye's passage has the following implication:

'Transparency Claim 7 (TC7):

When an object appears colored in a particular way to a person P in her visual experience, if P tries to turn his or her attention towards the intrinsic phenomenal character of her own color experience then she will not focus on (and she will not get aware of) any feature that she had not been aware of

already before.' (436).

She says:

'Transparency claim TC7 is in no conflict with the idea that experiences have intrinsic phenomenal character that we can attend to. But Tye seems to think that TC7 supports the contrary view. I suspect that he might be misled by the following implicit assumption:

Presupposition 2 (P2):

If it was possible to focus one's attention upon any intrinsic features of one's own experience then directing one's attention towards one's experience should involve getting aware of features one was not aware of before (when one's attention was not yet directed upon one's own experience).' (436-437).

It's not clear that TC7 entails WT, for it could be that TC7 were true but true because visual experience *always* involves awareness of a bifurcation of properties (mind-dependent/presenting and mind-independent/presented), so that when one 'tries to turn his or her attention towards the intrinsic phenomenal character' of his or her experience, she may do so, and still she will 'not focus on (and she will not get aware of) any feature that she had not been aware of already before'. But WT, if true, would explain why we should expect TC7 to also be true. Moreover, WT would explain why the following transparency claim, which Nida-Rumelin also endorses, might be true:

'Transparency Claim 6 (TC6):

When an object appears colored in a particular way to a person P in her visual experience, if P tries to turn his or her attention towards the intrinsic phenomenal character of her own color experience then she will still find herself attending to the color of the object'.

And she summarises: 'attending to the phenomenal character of one's color experience requires attending to the apparent color' (435). But if visual experience did *always* involve awareness of a bifurcation of properties (mind-dependent/presenting and mind-independent/presented) then it would be unclear

why TC6 should be true. For example, when viewing a painting we may attend to the paint strokes without also attending to the scene depicted. Nida-Rumelin has her own explanation for TC6. She proposes that attention to phenomenal properties is different because, unlike the painting, phenomenal properties are never objects of perception. Phenomenal properties, she says, are only ever objects of *introspective* awareness. Moreover, introspective awareness is not to be thought of on the model of perceptual awareness – this mistake, Nida-Rumelin argues, is the source of many errors on the topic of transparency⁸.

The above notwithstanding, we can also explain TC6 by appeal to WT, and it might be thought that WT has phenomenological plausibility in its own right. Moreover, although, as Nida-Rumelin correctly points out, neither TC6 nor TC7 is in ‘conflict with the idea that experiences have intrinsic phenomenal character that we can attend to’, as I shall explain below, it seems that WT *is* in conflict with that idea. In fact, WT is in conflict with any view on which phenomenal character is not identical with the objects of perception (i.e. any view other than Naïve Realism, and perhaps also Tye’s new (2015) version of Representationalism).

Let’s, for the moment, assume Weak Transparency:

When we introspect our visual experience we find but a single set of properties. We may think of these properties either as mind-dependent or as mind-independent. We do not find in visual experience duplicate sets of properties, regarding which we might think of the one set as mind-dependent (and presenting), while we think of the other set as mind-independent (and presented).

Assuming also that when we perceptually experience mind-independent objects we are introspectively aware of those mind-independent objects (which all supporters of any TE claim will endorse) then, on such occasions, any phenomenal features of our experience that might serve to present such objects to us, and which we must also be

⁸ However, see Martin (2006) for an argument that any Intentionalist (but not a Naïve Realist) will be committed to the perceptual model of introspection which Nida-Rumelin derides.

introspectively aware of (for we already established that phenomenal character must be introspectable – see the above discussion of Speaks' view), must be identical to (and therefore, in Soteriou's terms, sufficient for) those objects of perception. Otherwise there would be a bifurcation of introspectable properties of precisely the kind that WT vitiates against.

Though there does seem to be something correct about WT, it is very difficult to formulate it in such a way that it is not vulnerable to counter-examples. For example, as already noted, Block (1996), Kind (2003), and others have offered a number of examples of phenomenal features that, they say, do not introspectively seem to be mind-independent. An example from visual phenomenology is the blurriness of blurry vision (it's worth noting, though, that this could arguably also be seen as a counter-example to Nida-Rumelin's TC7). And these phenomenal features may be present even while other phenomenal features (a red table, for example), which do (or may) introspectively seem to be mind-independent, are also present in visual experience. In which case, we have a bifurcation of introspectable properties of precisely the kind that WT vitiates against.

Still it doesn't seem that this particular bifurcation of properties vindicates the idea that we find in visual experience duplicate sets of properties, regarding which we might think of the one set as mind-dependent (presenting), while we think of the other set as mind-independent (presented). But this depends upon exactly what is meant by the words 'duplicate' and 'presenting', and the trouble for Weak Transparency is that it is very difficult to make precise exactly what work these terms are meant to be doing. Indeed, without a precisification of these terms then even a local hallucination, which (in this case, we might suppose) is not liable to be mistaken by the subject for a veridical perceptual experience, within an otherwise veridical visual experience, might be thought a counter-example to Weak Transparency.

In this thesis I shall appeal to introspective observations in propounding an original argument in favour of Naïve Realism, over Intentionalism. The argument that I will present will not rely on any *general* transparency claim, of the kind that we have considered in this introduction. I will argue that Intentionalism implicates certain

specific phenomenal features in our visual experience that don't in fact seem to be present. And, as I shall argue, these are not features that Naïve Realism implies will be present in experience.

One can consider this to be an argument utilising a very specific kind of transparency claim. Because I claim that Intentionalism implicates the presence of specific phenomenal features that do not in fact appear to be realised, we can call the claim that these features are not realised a '*local* transparency claim'. This means that my argument will not face the two 'principle problems' that I have described, in relation to arguments that utilize what we can now call '*general* transparency claims'.

In contrast to WT, the transparency claims on which my arguments, in chapters 7-9, are based do not assume that all introspectable features of perceptual experience must be thought of as alike in their mind-dependence or mind-independence. So, unlike the above argument that utilizes WT, my argument is not vulnerable to the kinds of counter-examples posed by Block (1996) and Kind (2003) (the first principle problem with arguments utilizing TE). Moreover, this transparency claim does not make the contentious assumption that what presents itself to introspection is *manifestly* mind-independent (the second principle problem with arguments utilizing TE).

There is much ground to cover before we get to the local transparency argument. In chapter 1 I will outline and motivate three theory-neutral foundational theses. These theses are all related, and they deserve sufficient attention that to deal with them as and when they become relevant to my main argument would risk distracting from the thrust of the argument, so I will dedicate the first chapter to a consideration of them. As well as outlining the three theory neutral theses, I will also give some explanation of how those who accept different theories of perception spell out these theses differently.

In chapters 2 and 3 we shall consider how the Naïve Realist is to account for our perceptual experience of objects. This project is crucial to the plausibility of Naïve Realism, for as we have seen, it is common ground among the various

characterisations of TE that the properties that we are introspectively aware of are experienced *as* properties of objects. It may be a controversial question whether or not we experience such objects as mind-independent, but the claim that we experience properties as instantiated by objects, of some kind, seems to be on firmer ground.

Since the Naïve Realist describes perceptual experience as a relation between a subject, an object, and that object's properties, one might think that this is a requirement that the Naïve Realist can easily meet. However, as we shall see, matters are more complicated. In particular, though this simple description of the perceptual relation does build the object into the picture at the outset, it doesn't yet explain how the object manifests itself in perceptual experience *as such*. And given that objects consistently feature in the descriptions that we give of the phenomenology of our perceptual experience, objects must manifest themselves in perceptual experience *as such*. I will explain what exactly this means in due course. For now, suffice it to say, there appears to be no existing Naïve Realist explanation of this phenomenon, and that is something that I intend to rectify in chapters 2 and 3.

In chapter 4 I will offer some original phenomenological observations. I will argue that there is a pervasive feature of visual phenomenology that has so far gone unnoticed. This feature I will call 'phenomenal-composition'. I will offer a detailed elucidation of the phenomenon, the details of which, I hope, will at once aid the reader in finding it in her own experience, and also offer some insight into how the phenomenon has escaped detection until now.

Over the course of chapter 4, 5 and 6, I will argue that we perceptually experience spatial composition relations, and that phenomenal-composition is the manifestation, in perceptual experience, of those spatial composition relations. Crucially, I will also argue that we perceptually experience spatial composition relations *as* necessary relations. The argument presented in chapter 4 is a phenomenological argument in favour of these claims. Then, in chapters 5 and 6, I offer epistemological arguments in support of these same claims. In chapter 5 I argue that all alternative accounts of our knowledge of spatial composition relations are, at best, highly implausible. In chapter

6 I argue that all alternative accounts of our knowledge of the *necessity* of spatial composition relations are, at best, highly implausible.

In the final third of the thesis I use what has been established in earlier stages to argue in favour of Naïve Realism, over Intentionalism. This is where I utilise what I have called a 'local transparency claim'. The final third divides into three chapters, each of which address a different theory of perceptual content fixation. The first two such theories are Intentionalist accounts of content fixation, and the final one is a Naïve Realist account of content fixation. I will argue that only the Naïve Realist account of content fixation is capable of accommodating perceptual experience of the *necessity* of spatial composition in a way that is consistent with our visual phenomenology.

A caveat is in order. I believe that the argument that I offer, in chapters 4-9, *could* also be used to motivate the Sense Data Theory, but, as Sense Data Theory is today a very marginal position, I will not consider this application of the argument. The thesis should be read, then, as an argument in favour of Naïve Realism, on the assumption that Sense Data Theory is false. Before we begin, a point of terminology. Throughout most of the thesis I will refer to my favoured view as 'Naïve Realism', but, in chapters 2 and 3, I refer to this view instead as 'Relationalism'. These are existing names for the view that I favour, and my use of each of them in these different parts of the thesis is principled. In chapters 2 and 3 we shall be elaborating some structural features of the perceptual relation, so, in that context, it is useful to emphasise the relational nature of perceptual experience, on my favoured view of it. Then, in chapters 4-9, we shall learn how perceptual experience affords us with knowledge of essences, and I shall argue that this can only be so if my favoured view is correct. In this context the Naïveté of my favoured view comes into focus. So in this second context the name 'Naïve Realism' is the more apt.

Chapter 1 – Groundwork; The ‘Phenomenal Contents’ of Perceptual Experience.

In this chapter I will outline and motivate three theory-neutral foundational theses. These theses are all related, and they deserve sufficient attention that to deal with them as and when they become relevant to my main argument would risk distracting from the thrust of the argument, so I propose to dedicate this first chapter to a consideration of them. As well as outlining the three theory neutral theses, I will also give some explanation of how those who accept different theories of perception spell out these theses differently. This will allow me to introduce the different theories of perception that we’ll be considering in this thesis, in more detail.

1.1 Three Theory Neutral Theses Regarding Perceptual Content

1. **Content:** perceptual experience possesses content, where content is understood in a broad sense of the term, so as to be neutral between Intentionalist accounts and Naïve Realist accounts. Broadly, content is to be understood as the perceptual experience’s presenting the subject with a state of affairs.

In Intentionalist terms, possession of content means that the perceptual experience has accuracy conditions (Siegel: 2010, Tye: 2015, Horgan & Tienson, 2002: 524, 525, 526⁹). We can call intentional content ‘I-content’. It is in virtue of I-content that perceptual experiences, as the intentionalist understands them, are *about* the mind-independent objects, properties and relations (directly) perceived. In Naïve Realist terms, possession of content means that the perceptual experience is partially constituted by the objects, properties and relations perceived. We can call Naïve

⁹ Tye is a Representationalist and Horgan & Tienson are Phenomenal Intentionalists, so this construal of ‘intentional content’ is common ground among advocates of otherwise disparate versions of Intentionalism.

Realist content ‘NR-content’.

The Naïve Realist posits a primitive relation between the subject and the mind-independent objects, properties and relations (directly) perceived. According to the Naïve Realist, it is in virtue of this relation that perceptual experience is said to have NR-content. It is also consistent with Naïve Realism that perceptual experience has I-content. The view that I am arguing against is Intentionalism, which is the view that perceptual experience *only* has I-content (and no NR-content). Indeed, Soteriou (2013) is an example of a Naïve Realist who also attributes I-content to perceptual experience.

Miller (2017) thinks that accuracy conditions will characterize perceptual experience on any Naïve Realist account of it no less than it will on any Intentionalist account. But while there is a way of understanding the notion of ‘accuracy conditions’ on which Miller’s claim is true, this is not the sense of the notion at issue here. If Miller’s claim were true of the notion of ‘accuracy conditions’ that I am making use of, then clearly my two alternative definitions of ‘content’ would fail to distinguish Naïve Realism from Intentionalism.

We can get a view of the notion of ‘accuracy conditions’ regarding which Miller’s claim *is* true of by considering the following from Martin, which Miller quotes: ‘when we ask of the forgery of a Ming vase whether it is an accurate copy of the original, we do not assume that either the original or the copy ^[1]_[SEP] possesses a representational content. So in general, to ask of something whether it is accurate or not need not require it to be a representation or to have representational content, even if in some specific cases it is; it is simply to invite someone to match things’ (2010: 223)¹⁰.

It is easy to interpret the above passage as cleaving the notion of ‘accuracy conditions’ entirely from the notion of ‘representational content’ (or as I call it, I-content’). But if we interpret the passage in this way then we shall be left with the question of what distinguishes I-content (Martin’s ‘representational content’) from the mere

¹⁰ Martin, M.G.F. 2010. What’s in a Look? In *Perceiving the World*, ed. B. Nanay, 160– 225. Oxford: Oxford University Press.

possession of accuracy conditions, and we may be left with no option but to say that what distinguishes them is that I-content, unlike mere accuracy conditions, has logical structure, and this is exactly what Miller does say. This has the effect of assimilating all intentional content to specifically conceptual or propositional content. This position is problematic because most intentionalists believe that perceptual experience is possessed of content that is non-conceptual and non-propositional (See Tye: 1995, Crane: 2009, Cassam: 2014).

The copy of the Ming vase is only accurate or inaccurate relative to the intentions of the forger, or to the intentional states of whoever has been invited to match them. In this loose sense of the notion of 'accuracy conditions', absolutely anything can be possessed of accuracy conditions, provided that someone takes up the right attitude towards it. Certainly perceptual experiences, even on a purely Naïve Realist account of them¹¹, will be possessed of accuracy conditions in this loose sense. The fact that the subject of the experiences is inclined to make certain judgments on the basis of those experiences suffices to imbue her perceptual experiences with accuracy conditions in this loose sense. For example, on this loose notion of accuracy conditions, the fact that a subject is inclined to judge (and believe) that there is a cat on the mat, on the basis of a perceptual experience of a certain kind, suffices to imbue that type of perceptual experience with the accuracy conditions: 'the cat is on the mat'.

We can call the type of accuracy conditions that Martin is talking of '*relative accuracy conditions*'. What distinguishes Intentionalism from Naïve Realism has nothing to do with relative accuracy conditions. The distinction has its source in a difference in the way in which perceptual experiences are envisaged as providing judgments and beliefs with warrant or justification¹². What unifies NR-content and I-content, so that

¹¹ By 'purely', I mean to rule out hybrid views on which perceptual experience is possessed of both NR-content and I-content. See Soteriou (2013).

¹² This isn't the only thing that distinguishes Naive Realism from Intentionalism. The Naïve Realist also has a different view of phenomenal character to the Intentionalist. There may be other distinguishing factors too. The difference in their views on phenomenal character is intimately related to the difference in their conceptions of perceptual content. In thesis 3 we shall discuss the relationship

we may count both to be conceptions of perceptual *content*, is that they are both envisaged, within the respective theories, as being what affords perceptual experiences with the capacity to provide non-inferential warrant, or justification, for the judgments that are made on the basis of those perceptual experiences¹³.

Above I acknowledged that *relative* accuracy conditions may be imputed to perceptual experiences on a purely Naïve Realist account of them no less than on an Intentionalist account of them. This because the judgments that a subject is disposed to make on the basis of perceptual experience suffice to imbue those perceptual experiences with *relative* accuracy conditions. But given that the purpose of perceptual content is to confer upon perceptual experiences the capacity to provide non-inferential warrant, or justification, for the judgments that are made on the basis of those perceptual experiences, it should now be clear that relative accuracy conditions are not adequate to the task.

If content were conferred upon perceptual experiences by the judgments that they disposed subjects to make then that content would be incapable of justifying those very judgments. This is enough to show that the accuracy conditions attributable to perceptual experiences on a Naïve Realist account of them are not contents (at least in the sense that we shall use the term ‘contents’). So the two conceptions of content that I have described –NR and I – are genuine alternatives¹⁴. But what unifies these as different conceptions of one and the same thing is that both parties maintain that their perceptual contents will serve to present the subject with a state of affairs, and will thereby afford non-inferential warrant for perceptual judgments pertaining to that state of affairs.

Martin believes that the justificatory capacity of perceptual experience derives from the transparency of experience (TE):

between content and phenomenal character, and the corresponding difference in Intentionalist and Naïve Realist views on phenomenal character.

¹³ See Cassam (2014), who explicitly claims that this is the purpose of I-content.

¹⁴ While they are genuine alternatives, they are not mutually exclusive. It is possible to advocate a highbred view on which perceptual experiences involve both NR-content and I-content. Cf. Soteriou (2013).

‘The transparency considerations are relevant to an account of perceptual justification. It seems reasonable to us that we should come to believe that our environment is a certain way, given that our experience presents that environment as being that way’ (2002: 396).

If Martin is correct then there ought to be no qualms about saying that perceptual experience, on the Naïve Realist model, is capable of justifying our judgments about the mind-independent world¹⁵. As we saw in the Introduction, there is a simple argument that takes us straight from TE to Naïve Realism, and standard versions of Intentionalism (on which phenomenal character is distinct from the objects of perception) aren’t obviously able to accommodate TE. But why ought we to think that TE has anything to do with the justificatory role of perceptual experience?

The thought here seems to be that perceptual experience justifies judgments (or justifies subjects in making those judgments) by presenting to the subject the truth-makers for those judgments. Above I offered a definition of content that was neutral between Intentionalism and Naive Realism. This was: ‘content is to be understood as the perceptual experience’s presenting the subject with a state of affairs’. When the state of affairs thereby presented is the truth-maker for a judgment then the perceptual experience justifies that judgment, in virtue of its content. Martin is of the view that both Intentionalists and Naïve Realists can accommodate the justificatory role of perceptual experience¹⁶.

Now I turn to the question of what it is that determines the content of a perceptual experience. This is a very important question. Given that the content of a perceptual experience is meant to explain how that perceptual experience is able to justify the beliefs that we acquire on the basis of it, it is incumbent upon all parties to offer an

¹⁵ For another argument in favour of the claim that NR-content can afford a superior form of warrant to that which I-content is capable of affording, see Campbell (2002, 2014).

¹⁶ I suspect that there is in fact a tension in the ideas that, so far as the mind-independent world is concerned, perception has *only* I-content, and that perception has the capacity to justify judgments about the mind-independent world. The argument for this claim appeals to meta-theoretical considerations, and it requires a detailed examination of available views on the nature of justification, and so I will not take up this argument here.

account of how perceptual experiences come to have the contents that they do. Not only is this an inherently important question, but also the way in which this question is answered has implications for my argument in the final three chapters. Here I describe the theory neutral thesis, and I give a brief explanation of the competing specific accounts of it. In chapters 7-9 I will address each of these competing accounts in more detail.

2. **Content-Securing Mechanism:** Any notion of ‘content’ requires that there be some mechanism, consistent with that account, which explains how mind-independent properties and relations enter into the content of perceptual experience.

For the Naïve Realist, the content of the subject’s perceptual experience is determined by what properties she is perceptually related to. This, in turn, is determined by which properties fall in her visual field. So it is the environment (more specifically, that portion of the environment that falls in the subject’s visual field) that determines the content of the subject’s perceptual experience. The Naïve Realist’s explanation of how mind-independent objects and properties get into the content of perception is:

Acquaintance: There is an acquaintance relation between the subject, the mind-independent object, and its mind-independent properties. This might be polished as a relation of ‘conscious awareness’¹⁷.

As I shall explain later, this view is developed, by Campbell (2014), such that there is an acquaintance relation between the subject and all of the *properties* that fall in the visual field, and also an acquaintance relation between the subject and any *objects* that she visually attends to. What’s worth noting, at this point, is that the acquaintance relation itself does not bring with it any visual content. The

¹⁷ One might worry that this mechanism cannot be what secures for the Naïve Realist their idiosyncratic construal of content, for, as was mentioned in the introduction, Tye (2009) – an Intentionalist – now proposes an acquaintance relation. We shall discuss Tye’s recent view in chapter 6. In anticipation, Tye’s acquaintance relation is not a content fixing mechanism.

acquaintance relation just enables a certain portion of the subject's environment (the visual field) to determine a set of visual contents¹⁸.

For the intentionalist a perceptual state is individuated, or typed, independent of what is going on in the subject's immediate environment (and, more specifically, independent of what is going on in the subject's visual field). The content of a perceptual state is a matter of what accuracy conditions *that state* has. Since the accuracy conditions of a state are independent of what is actually in the subject's immediate environment, we are in need of an explanation of why a given type of state has the accuracy conditions that it does. This is a question that just doesn't arise for the Naïve Realist, and it is a question that the Intentionalist owes us an answer to.

One mechanism is ruled out at the outset by our motive for postulating content. The purpose of positing perceptual content is to explain how perception can provide non-inferential warrant, or justification, for the beliefs that we acquire on the basis of perception. This motive precludes accounts on which types of perceptual experiences derive their contents from the true judgments which they reliably give rise to. If the content of perception thus depended upon the contents of the judgments that it gave rise to, then the content of perception would not be capable of providing warrant for such judgments¹⁹.

I will consider two types of mechanism by which mind-independent properties and relations might enter into the intentional content of perceptual states:

¹⁸ This is why Naïve Realists are generally committed to disjunctivism (See Soteriou, 2016).

¹⁹ However, see Lewis (1980) for an apparent endorsement of this view. Chalmers' (2006) also, in some respects, resembles this view. He advances a causal view of content fixation, but one on which either the individual, or society, imposes causally determined content upon perceptual experiences. Yet, at the same time, Chalmers seems to want to say that the individual, or society, are responding to objective standards of rationality in doing so. But the whole point of attributing content to perceptual experience is to engender such normative standards for responding to those perceptual experiences. There is a definite whiff of circularity here. For our purposes we can sidestep this issue and simply class Chalmers' view as a Causal View of content fixation, so the argument in chapter 7 will apply to his view.

1. Causal Accounts: A causal relation between types of perceptual states and mind-independent features of the environment gets those features into the content of the perceptual states.

And

2. Benefit-Based Accounts: the content of some representation is whatever feature of the distal environment happens to be responsible for the fact that the representation *benefits* the organism.

The abstract idea, expressed in 2., is often described using architectures involving representations, representation producers, and representation consumers. Those Benefit-Based accounts that utilize architectures of this form are called 'Consumer Semantics Accounts'. When described in this way, the mechanism can be given the following, more concrete, explication:

(2.1) Consumer Semantics: the relevant sub-personal representation consumer determines the content of a representation. There is a mapping relation that must obtain between representations and the environment if those representations are to enable the relevant representation consumer to perform its proper function.

Some theorists advocate narrow mental content. But even among these authors there is broad consensus that there need additionally be wide mental content. Thompson (2010), Chalmers (2005, 2006), Horgan and Tienson (2002: 528-529), Horgan et al. (2004), Shoemaker (2000), and Fodor (1987), all believe that there is wide content in addition to narrow content²⁰. I am claiming that, even if we take narrow mental

²⁰ Some caveats are in order. When Horgan et al. talk of wide content, they talk of *thoughts* concerning particulars and natural kinds. What's more, they say, of perceptual experience of properties and relations, that 'such mental reference is wholly constituted phenomenologically' (2004, 304), where phenomenology is assumed to be narrow. Nonetheless, they offer no account of how such wholly phenomenologically constituted reference to mind-independent properties is possible. So it may be (indeed, I think it will be) that Horgan et al. require wide perceptual content after all. Regarding Fodor, what is needed is a caveat running in the opposite direction. Fodor's (1987) narrow content is more

content for granted, we ought to expect an account of how perceptual states get their wide content. Fodor (1987, 1990), Thompson (2010), and Chalmers (2006), evidently agree, as they all explicitly endorse a causal account of wide content.

The Composition Objection – the focus of chapters 7-9 – applies differently to each of these views of content-fixation, so each of these views will receive a bespoke treatment, in separate chapters. I now turn to the final thesis, which pertains specifically to the *phenomenal* spatial contents of perceptual experience. This final thesis will be crucial when it comes to arguing that Naïve Realism, but not Intentionalism, can accommodate perceptual experience of the necessity of spatial composition.

3. **Phenomenal Spatial Contents:** A spatial property or relation enters into *phenomenal* content if and only if that property or relation enters into the content of a perceptual state, *and* there is an element of phenomenal character that uniquely corresponds to the perceptual content.

This relation between contents and corresponding elements of phenomenal character is crucial to my argument in chapters 2 & 3, and again in chapters 7-9. In chapters 7-9 I will argue that, given Thesis 3 above, and a certain phenomenological observation to be described in chapter 4, neither Causal views of content fixation nor Benefit-Based accounts of content fixation can accommodate perceptual experience of the necessity of spatial composition. In contrast, I will argue in chapter 9 that Naïve Realism *can* accommodate perceptual experience of the necessity of spatial composition, in a way that is consistent with both Thesis 3 above, and with the phenomenological observation to be described in chapter 4.

Now I turn to the nature of the ‘correspondence’ relation between content and

anaemic than that of the others (Fodor’s narrow content is not truth functional). Finally, there are those - Farkas (2008) – who believe that there is *only* narrow content. But Farkas is the exception in this regard, and she faces a similar problem to Horgan et al., because she offers no account of how phenomenologically constituted content is able to serve as a function from a context to truth conditions involving mind-independent properties.

element of phenomenal character. A type of element of phenomenal character corresponds to a type of content only if there is a dependence relation between the content type and the phenomenal type. It could be that the phenomenal type asymmetrically depends upon the content, or the content type asymmetrically depends upon the phenomenal type. Alternatively, the phenomenal type and the content type might be identical, or the phenomenal type and content type might both have a common dependence upon some third item (a representational vehicle).

Finally, the dependence relation must be *unique*. If the phenomenal type is associated (in the way described above) with two contents then that phenomenal type does not suffice to get either of those contents into a perceptual experience; at most it would get a content indeterminate between them into the content of perceptual experience.

A.D. Smith, in his (2002) *The Problem of Perception*, fails to recognize the point at issue here. For A.D. Smith it is the 'sensuous' nature of perception that generates the argument from illusion, against direct realism. He also appeals to the 'sensuous' nature of our experiences to motivate the primary quality/secondary quality distinction. Secondary qualities, he claims, are those that are responsible for the 'sensuous' features of perception. He calls these sensuous features 'sensations'. Assuming that sensations are, for A.D. Smith, co-extensive with phenomenal character, the implication is that there are no elements of phenomenal character that correspond to primary qualities.

It's not entirely clear, however, that A.D. Smith does envisage sensations to be co-extensive with phenomenal character, for he denies that sensations are intentional, and he also says that his goal is to:

'Work out an adequate analysis of perceptual consciousness that can do justice to its intentionality, its phenomenological world directedness, while construing the sensory qualities that are present in such consciousness as intrinsic states of the experience itself' (2002, 58).

So, according to Smith, 'world directedness' is distinct from 'sensation', yet the 'world

directedness' is distinctly 'phenomenological'. But if something can show up in phenomenology without being a sensation, what purpose do sensations serve? Indeed, what is the difference between phenomenology and sensation? Assuming that a property is 'sensuous' if and only if perceptual experience of that property is uniquely associated with some phenomenal type, A.D. Smith's distinction between primary qualities and secondary qualities cannot be right, because perceptual experiences of spatial properties clearly are uniquely associated with specific phenomenal types: There's a difference, phenomenologically, between seeing a *square* patch of green and seeing a *triangular* patch of green. Indeed, plausibly, 'It is beyond doubt that one can phenomenally represent squares as such' (Bayne: 2009, 401).

Still, the demand that there be *some* dependence relation between content types and phenomenal types does not yet prejudice us against any extant views on the metaphysics of phenomenal character. Indeed, it is consistent with this demand that:

- a) Each visual, spatial, phenomenal-type *uniquely* depends on one spatial perceptual experience content-type.
Or
- b) Each spatial perceptual experience content-type *uniquely* depends on one visual, spatial, phenomenal-type.
Or
- c) Each visual, spatial, phenomenal-type is identical with a spatial perceptual experience content-type.
Or
- d) Each visual, spatial, phenomenal-type *uniquely* depends on one spatial perceptual representation vehicle-type, and each spatial perceptual experience content-type also *uniquely* depends on one spatial perceptual representation vehicle-type (from the very same set of spatial perceptual representation vehicle-types).

1.2 Summary

The argument that follows will be constructed around three theses:

1. **Content:** perceptual experience possesses content, where content is understood in a broad sense of the term, so as to be neutral between Intentionalist accounts and Naïve Realist accounts. Broadly, content is to be understood as the perceptual experience's presenting the subject with a state of affairs.
2. **Content-Securing Mechanism:** Any notion of 'content' requires that there be some mechanism, consistent with that account, which explains how mind-independent properties and relations enter into the content of perceptual experience.
3. **Phenomenal Spatial Contents:** A spatial property enters into *phenomenal* content if and only if that property or relation enters into the content of a perceptual state, *and* there is an element of phenomenal character that uniquely corresponds to the perceptual content.

Once spelt out, none to these theses are controversial. I think that the only reason that some authors²¹ have said things that appear to be inconsistent with these theses must be because these theses are rarely (if ever) explicitly stated. With this done, we can now move onto the argument proper.

²¹ A.D. Smith (2002).

Chapter 2 – Naïve Realism and Perceptual Experience of Objects.

For reasons already explained, in this chapter, and the next, I shall refer to Naïve Realism as ‘Relationalism’. In section 1 I describe an important, and fundamental question facing the Relationalist, and I explain how the beginnings of an answer to this question are implicit in Campbell’s contribution to Berkeley’s Puzzle (2014). Perceptual experience, on Relationalist models of it, is usually thought to consist in a relation between a subject and an object, and not just between a subject and some collection of properties. There is, however, an important question about what justifies the Relationalist in describing the perceptual relation in this way. Campbell (2014) draws our attention to a distinction – the selection-access distinction - in the light of which we can see why the perceptual relation ought to be thought of as involving objects, and not just their properties.

In section 2 I consider how, on the present proposal, objects might present themselves in perceptual experience *as such*. On my interpretation of Campbell’s view, selection serves to present the object that is selected as identical throughout the period during which it is selected. I explain how, in these cases, selection also serves to justify the subject’s judgements about the object’s identity.

2.1: Why a Perceptual Relation with *Objects*?

Relationalists think that perceptual experience is a relation between a subject, an object, and the object’s properties. Relationalism, as I understand it, claims that (at least part of) the phenomenal character of perceptual experience is to be identified with the perceived properties of objects. The fact that the perceived properties are identical with (at least part of) the phenomenal character of one’s perceptual experience is what provides the rationale for the claim that, in perceptual experience, one is perceptually related to the object’s properties. But what justifies the Relationalist’s further claim that, in perceptual experience, the subject is related to

the *object* that possesses those perceived properties? Another way of putting the question is this: What justifies the Relationalist in denying that the subject is only perceptually related to properties (and not to objects)?

One might object to the demand for an answer to this question in the following way. As Armstrong (1989) puts it, properties are just the *ways* that things can be, and objects (or substances) are just the *things* that are those ways. You cannot have a property instance or trope – a token *way* of being – without having an object – a *thing* that is that way. So it follows, the thought goes, that in perceiving a property – a *way* of being – one necessarily perceives the *thing*, which is that way.

Perhaps there is a sense in which merely perceiving an object's properties suffices to also perceptually relate one to the object. But this is not the sense that Relationalists have in mind when they describe perceptual experience as a relation between a subject, an object, and its properties. If it were then it would arguably be redundant to mention the object, since, on this view, the object *only* enters into the relation because its properties do. So nothing substantive is added by saying that the subject is also perceptually related to the object. I think that when Relationalists describe the relation as involving objects they mean to say more: they mean to say that the object features in the *phenomenal* content of perceptual experience. This means that the object manifests in the phenomenology of perceptual experience as an element *distinct* from its properties²².

One reason in particular for thinking that the Relationalist ought to say that the object features in the *phenomenal* content of perceptual experience is that, as Martin (2002) points out (and as I mentioned in chapter 1), Relationalists tend to think that perceptual experiences afford us with non-inferential warrant for judgements in virtue of the way those experiences *seem* to subjects undergoing them. The 'seeming' in question here appears to be a *phenomenal* 'seeming'. So if perceptual experiences

²² This should not be taken as implying that *whenever* we perceptually experience an object's properties then the object features in the phenomenal content of perceptual experience.

are to provide us with non-inferential warrant for judgements concerning objects then objects had better feature in the phenomenal content of perceptual experience.

As Martin (2002) notes, it is partially definitive of perceptual experience (as opposed to, for example, imagination) that perceptual experiences dispose subjects towards believing their contents. If one perceptually experiences a cat on a mat then one is disposed to believe that there is a cat on the mat. If one imagines a cat on the mat then one is not thereby disposed to believe that there is a cat on the mat.

The Intentionalist says that when one perceives something one bears the 'perceiving' attitude to the relevant content, and this attitude is such as to dispose the subject to believe the content. As Martin (2002) describes the Intentionalist's view, this disposition is actually responsible for the fact that in perceptually experiencing something it thereby *seems* to the subject as though her environment is the way described by the relevant content (this, Martin claims, is how the Intentionalist accounts for the Transparency of Experience)²³. But the Relationalist is unlikely to want to endorse this order of explanation. The Relationalist is more likely to want to say that the disposition is explained by the fact that, in perceptually experiencing some state of affairs, it thereby *seems* to the subject as though that state of affairs obtains.

The initial problem with this is that there is a sense in which it is in tension with the phenomenology:

“The phenomenology of object experience seems to present us directly with objects, but it does not seem to acquaint us with their intrinsic nature in a sense over and above acquainting us with their colors, shapes, and so on. If it did, then the phenomenology of object experience would be quite different from what it is: experiences of different tennis balls would typically have quite

²³ Martin may specifically have Representationalism in mind here. A Phenomenal Intentionalist might wish to reverse the order of explanation, so that the disposition is explained by the fact that, in perceptually experiencing some state of affairs, it thereby *seems* to the subject as though that state of affairs obtains. This is the order of explanation endorsed by Relationalists.

different phenomenal characters, for example. But the experience of objects does not seem to be this way.' (Chalmers, 2006, 108)

He also adds that 'The phenomenology of perception does not seem to reveal the intrinsic haecceitistic natures of objects' (Chalmers, 2006, 109).

But if 'the phenomenology of object experience does not seem to acquaint us with their intrinsic nature in a sense over and above acquainting us with their colors, shapes, and so on', in what sense does 'the *phenomenology* of object experience seem to present us directly with objects' (italics added)? And this question is all the more acute when we keep in mind that perceptual experience of objects ought to provide us with non-inferential warrant for our judgments about them. For, as I argued in chapter 1, this requires that objects enter into the content of perceptual experience (this was thesis 1). And moreover, I there argued that if something features in the content of perceptual *experience* then there must be an element of phenomenal character that uniquely corresponds to the relevant content (this was thesis 3).

This means that if objects, as well as properties, feature in the content of perceptual experience then there must be an element of phenomenal character that uniquely corresponds to the object (and not to any of its empirical properties, which we also perceive, and which we're seeking to distinguish it from). From this vantage point it can seem as though we are faced with an impossible task. The Relationalist believes that properties enter into the content of perceptual experience by being identical with some element of phenomenal character. But the only property that distinguishes the object from the properties that happen to define it at a given moment is the object's *identity*. And identity, being a metaphysical property, just doesn't intuitively seem to be something that could be an element of the phenomenal character of one's perceptual experience.

One reason for scepticism here is that it is usually assumed that, if we are to perceive some property, our visual systems must be causally responsive to that property. But there is a sense in which the causal powers of an object seem to be exclusively determined by that object's empirical properties. After all, once one knows all of some

object's empirical properties, in theory, one can know exactly how that object will behave. And Chalmers' phenomenological observation (above) seems to offer some support for this view.

It becomes possible for the Relationalist to offer an account of how objects show up in the phenomenology of perceptual experience when we recognise that objects and property types manifest themselves in perceptual experience in fundamentally different ways. The following is (in a somewhat more explicit form than Campbell ever arrives at) Campbell's answer to the question of how objects manifest themselves in perceptual experience:

Objects manifest themselves, in visual experience, in the structure of our conscious visual attention to the properties in our environment.

There is an element of phenomenal character here, but the element is a structural feature of our experience. The element of phenomenal character is fundamentally unlike those that correspond to the empirical properties that we perceive (shapes, colours, etc.), for it is a higher order feature of our experience that depends upon our experience of those empirical properties. Moreover, the significance of this element – as the presentation of an object *as such* – is one that only becomes apparent when one considers its temporal features. I will say more on the temporality of this element of phenomenal character in section 2. First, I will describe the initial empirical basis for Campbell's claim and I will consider the most immediate philosophical questions that this empirical research gives rise to.

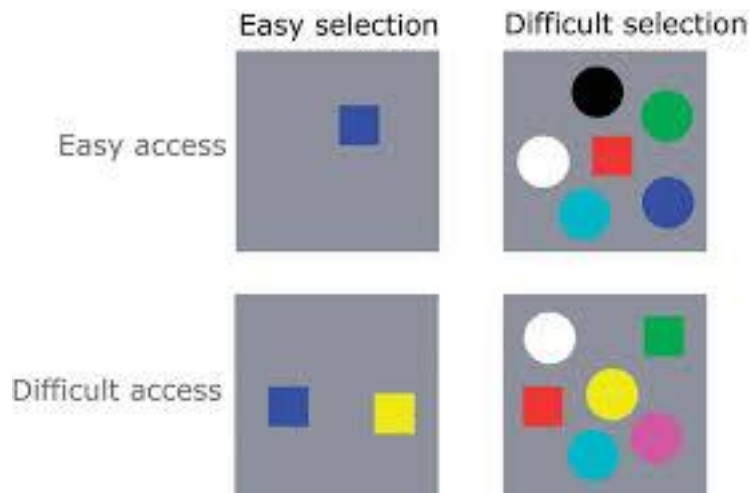
The above statement is Campbell's answer in its most abstract form. A more concrete exposition of Campbell's answer requires that we introduce a distinction from the literature on visual attention. Huang and Pashler (2007) propose a distinction between two types of visual attention: selection and access. The task of the rest of this section will be to provide an explanation of the selection-access distinction, and of the role that Campbell sees it as playing.

‘In visually attending to a scene, one dimension of your experience has to do with the characteristics of the objects that you would report them to have, act with respect to, or report yourself as experiencing. But another, more fundamental dimension of visual experience has to do with how you grab the object in the first place; how, in vision, you snatch it out from the rest of the visual array as something on which you are going to focus’ (Campbell: 2014, 51).

The distinction that Campbell is talking about is a distinction between two different types of visual attention: selection and access. The latter, ‘more fundamental’ dimension of visual experience is selection. It is Campbell’s contention that selection is the means by which objects enter into the perceptual relation. As Campbell explains, the distinction is best explained in terms of two dimensions along which the difficulty of certain kinds of tasks can be varied.

A task that serves to elucidate the distinction is one in which subjects must report the colours of squares in a display. In order to perform this task subjects must select each square, one at a time, ignoring any non-squares, and access and report the colour of each one. The task can be made more difficult in either of two ways. The first way in which the task can be made more difficult is by introducing more distractors – more non-squares -, which must be ignored. The second way in which the task can be made more difficult is by introducing more squares the colour of which must be accessed. The first way of increasing the difficulty of the task increases the difficulty of selection, and the second way of increasing the difficulty of the task increases the difficulty of access.

Below is a diagram, first printed by Huang and Pashler (2007), and reprinted by Campbell and Cassam (2014). The diagram describes the difference between selection and access, as dimensions of difficulty in a task.



1.

The diagram describes the difference between selection and access, as dimensions of difficulty in a task. The diagram was first printed by Huang and Pashler (2007), and subsequently reprinted by Campbell and Cassam (2014).


The explanation that Huang and Pashler offer of the distinction, at the level of neural processing, assumes an underlying architecture of Treismanian feature maps and Boolean maps. The idea is that there is a feature map for each type of feature (or property dimension) – one for colours, one for shapes, one for motion, one for orientations, and so on. Each feature map plots the locations of all of those properties in the subject’s visual field, of the relevant kind. Selection serves to initiate the creation of a Boolean map, on which features of different kinds (or properties belonging to different dimensions) – say, a shape, a colour, and an orientation – may be combined.

As can be seen on the diagram below, selection is under ‘command from top-down control’²⁴. The subject can generate a Boolean map in either of two ways. The subject can use a feature value – say, yellow, or square – to select the regions of the visual field occupied by that type of property, and features of other kinds (from other property dimensions) will be automatically added to the Boolean map. For example, if the subject selects the regions of her visual field at which yellow is to be found, and there is only one yellow object in her visual field, then the shape of this object will be automatically encoded on the Boolean map. Alternatively, the subject can use a set of

²⁴ In fact, Huang & Pashler claim that ‘all top-down control is ultimately attributable to selection’ (2007: pdf page 19)

locations to generate a Boolean map detailing the features to be found at those locations. Once the subject has done this then she may go on to access the properties that are represented on the Boolean map.

The Boolean map is subject to some stark limitations. It may only encode one value for each property dimension. So, there may be only one colour, one shape, and one orientation encoded on each Boolean map. If the region selected contains two different colours then the Boolean map (assuming that it is veridical) will be neutral regarding colour. Likewise, the Boolean map will encode the shape of the entire region selected, but not of any sub-regions. If one selects a region containing two colours – perhaps an object with yellow parts and blue parts -, and one wishes to access both of these colours, then one will have to first generate two further Boolean maps - one pertaining only to the yellow sub-region, and the other pertaining only to the blue sub-region. Then one will be in a position to access these colours. Below is a diagram of the architecture:



This diagram, first printed by Huang and Pashler (2007), and reprinted by Campbell and Cassam (2014), illustrates the information processing that is hypothesised to underlie selection and access.

Huang and Pashler (2007) assume that only what the subject visually accesses is conscious (i.e. a part of her visual experience), but Campbell points out that this interpretation is entirely optional. There is no reason to assume that visual experience is exhausted by what the subject is in a position to report on, or act with respect to²⁵. Our common sense view of visual experience is that we are simultaneously presented with a rich manifold, or array of entities, in our immediate environment. This manifold, or array, which we standardly call the 'visual field', is defined by a spatial region, relative to the subject's eyes, the contents of which are *simultaneously* visible to the subject. On the common sense view, the visual field simultaneously presents us with a range of entities that we may then go on to single out for closer inspection.

There are many examples that support the common sense view. One may be entirely absorbed in conversation, so that one does not notice the chiming of the church bells. It may only be when the bells stop chiming that one becomes cognitively aware both of their absence and of their previous presence. But the simplest explanation of this is that, though one was not *attending* to the sound at the time, one did previously experience the chiming of the bells. On this view, the onset of silence simply causes one to attend to that prior experience of chiming.

It is not just phenomenology that supports the common sense view. We ordinarily believe that we can *choose* which of the entities in our visual field we go on to visually attend to. Huang & Pashler seem to want to honour this conviction, because they say that selection is under 'command from top-down control' (see the diagram above). But if we are only conscious of what we have actually attended to then these commands from top-down are completely blind. Why do we select one region of the visual field rather than another, if we have absolutely no idea what occupies any of them?

²⁵ In assuming that visual experience is exhausted by what the subject is in a position to report on, or act with respect to, Huang and Pashler are agreeing with the standard interpretation of the phenomena of change blindness and attention blindness. In opposing Huang and Pashler on this point, Campbell is also offering an alternative interpretation of these phenomena - one that reconciles our common sense view of visual experience with the phenomena.

Some authors (Cohen et al.: 2012) have denied the common sense view. These authors say that it merely seems as though we are simultaneously aware of entities stretching right across the visual field because we can turn our gaze towards them so quickly. Some have called the common sense view the 'grand illusion'. But it is generally accepted that the imputation to subjects of systematic error in any area is to be avoided, if at all possible.

In defence of the common sense view Campbell suggests that in order to report on some feature of the visible environment, or to act with respect to it, closer inspection is required. But since the content of the visual field is independent of, and determined prior to, what the subject actually goes on to inspect, it is a mistake to assume that only what is reportable is conscious (i.e. part of the subject's visual experience).

Campbell offers the following example in support of this claim:

'A tiger padding through the veldt may be able to distinguish its prey from the foliage because of the colour of its target; but that does not mean that the tiger has any interest in the colour of things; it may not even be able to attend to the colour of the object. All it cares about is the object itself. The tiger may be incapable of attending to the colour of the object, even though it uses the colour of the thing to select the object from its background' (2014, 61).

Campbell's view is that the event that he describes above can be also described at the level of information processing in the terms of Huang & Pashler's theory. But though the event can be described at either the personal (or creature) level or at the information processing level, it is but a single event that is occurring.

The common sense view, on which the visual field presents us with a range of entities that we may then go on to single out for closer inspection, can be given a very precise formulation in the terms of the selection-access distinction. And in formulating the common sense view in these terms we can see how objects enter into the perceptual relation. In these terms, to visually experience some property is for that property to

be available (because it falls within the visual field) for us to use to *select* an object as figure from ground. When we thus *select* an object as figure from ground, we are then able to inspect, or *access*, that object's properties. So objects enter the perceptual relation just when we *select* them, and visual experience outstrips what is actually *accessed*²⁶.

The difference between Huang and Pashler's interpretation and Campbell's interpretation can be summed up as follows: For Campbell, the selection-access distinction is one that is to be found *within* perceptual experience, and it has a direct role to play in characterising the nature of our perceptual experience. For Huang and Pashler the distinction is only to be found at the sub-personal level, and it only has a role to play in determining which properties in our environment are perceived and which of *those* properties thus characterise our perceptual experience.

Campbell offers some examples that he hopes will evoke intuitions sympathetic with his interpretation of the distinction. He suggests that if we are using some property – say, colour – to select an object as figure from ground then we must be conscious of that property, even if we are *incapable* of going on to access that same property. On the first point – that if we are using some property to select an object as figure from ground then we must be conscious of that property –, Campbell says the following:

'Let's take a relatively simple kind of example: the use of tests for colour vision in which all that separates a particular figure from its background is the hue. So, for example, you may be presented with an array of variously coloured blobs of varying luminance in which all that there is to systematically separate a figure 5 from its background are the colours of the various blobs (Ishihara colour test plates). In such a case, it is hard to imagine how you could see the 5 without having conscious experience of the various colours involved... if the 5 is visibly there, as a 5 can be present in ordinary vision, then we cannot

²⁶ One might be surprised to hear me talking of selection as something that *we*, rather than our visual system's, do. Later I will explain that selection is usually an action, under the control of the subject.

imagine how that could be unless you had phenomenal awareness of the colour' (2014, 57).

On the second point – that it does not follow from the fact that we are using the property to select the object that we are therefore capable of accessing that property -, Campbell says the following:

'It is not difficult to find illustrations... of the possibility of selection on the basis of a property without access to that property. Human children have colour vision in place long before they have anything in the way of a colour vocabulary; it is entirely possible that a child a few months old could see the 5 in the kind of display I am describing, without having any ability to give a verbal report of its colour... A mere ability to see the 5 against the background does not imply that one has a capacity for colour induction. Nor, to take some further obvious examples, does it imply that one has a capacity for colour matching, or a capacity for ordering objects by their colours'. (2014, 59)

Campbell concludes that we ought to:

'Regard a property's figuring in awareness as a matter of that property being available for use as a basis for selection of an object or region (rather than as a matter of that property being accessed in vision)' (2014, 60)²⁷.

Finally, Campbell reaches an account of our perceptual experience of objects. In his words:

²⁷ Why does Campbell not say that properties only figure in experience when they are actually used to select an object as figure from ground? One reason is that we can only select one object at a time (more on this in the next chapter), so this would be inconsistent with the common sense view of our phenomenology, on which we are simultaneously presented with properties stretching across our visual fields. Another reason is that selection is usually an action (more on this later in the present chapter), and if we were not already conscious of properties before we used them to select objects then it's difficult to see how we could choose to select the relevant object.

‘The correct way to formulate a relational account of perceptual experience is to think of the relation as holding between a thinker and an array of visible properties at various locations, available for use in the selection of objects as figure from ground. Objects figure in sensory experience only when selected as figure from ground, ready to have their further characteristics accessed’ (2014, 64-65).

Campbell’s (2002) description of the perceptual relation is that it is a three-place relation between the subject, the object, and the subject’s point of view. Campbell now construes the perceptual experience that one has of the array of visible properties, which constitute the visual field, as a two-place relation between the subject and those properties. He then describes the properties that we use to select the object as figure from ground as the ‘mode of presentation’ of the object. The implication seems to be that we ought to construe the perceptual experience that one has of the object as a three-place relation between the subject, the object, and the subject’s point of view, where the subject’s point of view is now analysed as the properties that the subject uses to select the object as figure from ground (that is, the ‘mode of presentation’ of the object).

But why think that we select objects, rather than just regions of space? For selecting the region of space that the object happens to occupy would afford the subject with access to many of the same properties. And though the objects do instantiate some properties that the regions of space do not – the property of being an object of a given kind, for example – the Relationalist is unlikely to want to say that we perceptually experience such properties. So what could distinguish between these two possibilities?

Campbell could be clearer on this point, but the answer seems to be as follows: Visual access to properties is not instantaneous, and if the object is continuously moving then one will have to track the object if one is to visually access its properties. Clearly, if the subject is actively tracking the object, and this enables the subject to access the properties that she is perceiving, then the object needs to figure in the description that we give of the perceptual relation. And matters are not altered substantially just

because it so happens that, in a given case, the object is not in motion. It seems sufficient just that the subject would track the object were the object to move. So we can say that if the subject is disposed to track the object then the object figures in the perceptual relation. Usually we are disposed to track objects, and at least in those cases it is an object, rather than a region of space, that we are selecting.

Here is Campbell's characterisation of the situation:

'There may be nothing systematic to be said in terms of properties about how the subject is keeping track of the thing over time; it may be changing as it moves, and perhaps the only concise systematic thing to say about the principles governing the subject's selection of the thing over time is that the selection is sustained by it being the same object over time' (2014, 70).

This, then, provides us with an answer to the question why we ought to think that the perceptual relation involves the object, and not just its properties. However, there are other complications with Campbell's proposal. As mentioned above, Martin (2002) notes that it is partially definitive of perceptual experience (as opposed to, for example, imagination) that perceptual experiences dispose subjects towards believing their contents. The Relationalist wants to say that this disposition is explained by the fact that, in perceptually experiencing some state of affairs, it thereby *seems* to the subject as though that state of affairs obtains. In the present case this means that we need to explain how it is that in being perceptually related to an object it thereby *seems* to the subject as though she is perceptually related to an object.

2.2 - The Presentation of Objects *as Such*.

It is a point that has been widely observed that syllogisms like the following are not formally valid, though the conclusion does necessarily follow from the premises:

P1: The morning star is a planet in our solar system.

P2: All planets in our solar system orbit the sun.

C: The evening star orbits the sun.

The conclusion necessarily follows from the premises because 'the morning star' and 'the evening star' refer to the same object - Venus. But the syllogism is not formally valid because one may understand, and accept, the premises, and still coherently doubt the conclusion, if one does not know that 'the morning star' and 'the evening star' refer to the same object. This means that, although the premises do necessitate the conclusion, acceptance of the premises alone does not justify acceptance of the conclusion. Similarly, if one is visually related to a single object over some period of time then this does necessitate the existence of that object throughout the period of time (you can't be related to something that doesn't exist). But unless it also thereby *seems* to one as though one has been visually related to the same object throughout that period of time then one will not be disposed to believe that one has been visually related to the same object throughout that period of time.

In section 1 I argued that objects enter into the content of perceptual experience, and in making this argument I appealed to facts about the subject's tracking dispositions, and how these relate to what is in fact in the subject's environment. But, as has now become clear, that the subject does in fact track a single object over time, or that they are in fact disposed to do so, is not yet sufficient to show that the object features in the content of the perceptual experience. For this to be so it must also thereby *seem* to the subject as though she has tracked, or is disposed to track, a single thing – an object – over that period of time.

If we can explain how the identity of objects manifests itself in perceptual experience then this will also serve to explain how the identities of property instances (or tropes) manifest themselves in perceptual experience. Moreover, this seems like the correct order of explanation, since, save identification by location, the only non-demonstrative way we have of referring to a property instance (or trope) is by identifying the object that it belongs to, and then specifying the property type. My proposal is that, through selection, the object manifests itself *as such*, and it manifests itself *as identical* throughout the period over which it is selected.

I wish to propose an explanation of how objects manifest themselves as identical over time, which leans on the object-directed nature of visual attention. There is a wealth of empirical data supporting the claim that attention is *inherently* object-directed. As Scholl puts it:

‘The units of attention are often various kinds of visual objects. That this is true seems undeniable in the face of converging evidence from so many psychophysical and neuropsychological experiments’ (2001: 39).

Bringing the literature on object-directed attention together with Huang & Pashler’s model, we can say that whenever we make a selection there is always a set of tracking dispositions that are thereby mobilised, or engaged (I will come to the empirical evidence in support of this claim, in chapter 3, once we have a clear view of the proposal). These tracking dispositions have assumptions about the identity conditions of objects built into them - for example, that objects may survive changes in their location properties, provided that these changes describe a continuous path through space. These dispositions are the result of bottom-up processes in the visual system. The subject is able to override the dispositions, as she does when she chooses to visually attend to a different object, but even when the subject does override them, the dispositions are still inherent to the operation of selection that is being executed.

In addition to the bottom-up processes that enable us to track selected objects as they move there are also bottom-up processes that continually serve to orient us, and to orient the parameters that we use for selection, towards objects in the first place. I will turn to the empirical literature on object-directed attention in chapter 3, but for now, the following quote should help to make clear just how robust this literature is:

‘In few other areas of this young field have so many areas of study converged on so many similar ideas, and as such the research on this topic might be viewed as an emerging ‘case study’ in cognitive science’ (Scholl 2001: 39).

Drawing on this rich empirical literature, my proposal is that we can explain the phenomenology of being related to objects if we just assume that the subject is

generally aware of whenever a *new* selection has been made. Huang & Pashler (2007) explicitly assume that selection is something that we exert some control over (as is made clear by the inclusion of a reference to ‘top-down’ control, in the diagram above), and it is a commonplace assumption that we have non-observational knowledge of our own actions. I can choose to focus first on the coffee mug (providing me with access to the properties of the coffee mug, but also causing the rest of the items on the coffee table to melt into the background), now on the wicker basket (providing me with access to the properties of the wicker basket, but causing the coffee mug to melt into the background, along with the rest).

The appearance of a new object in the visual field will ‘capture’ our attention, where ‘capture’ is a purely bottom-up and involuntary process (Hillstrom and Yantis 1994, Yantis and Hillstrom 1994). But if the subject wasn’t aware of when such attention ‘captures’ have occurred then the subject would not be in a position to orient her attention back towards the object that she is concerned with. Yet subjects will immediately orient their attention back, if the new object does not bear on their present project.

Given the bottom-up, object-tracking, processes that I have alluded to, we can be confident that, generally speaking, whenever a selection has been made at t_1 , and by t_5 no further selection has been made, a single object has been selected throughout the period, t_1 to t_5 . We also now have reason to believe that, in such cases, the subject will be aware that no new selection has been made during this period. The implication is that the subject knows the difference between the case in which, from t_1 to t_5 , she selects and tracks a single object and her experience changes because that object’s location or other properties change, and a case in which she switches her attention to a new object at t_3 , and her experience changes because she is attending to a different object.

The present proposal, then, is that through selection the object manifests itself *as such*, in that, if we noticed how our experience unfolds during a single selection, we would conclude that the course of our experience was due to the (changing) nature of a single object. Those experiences that unfold when we do not make any new

selection, and when we allow our tracking dispositions to work freely, are what constitute the phenomenology of being perceptually related to a single object. This is the sense in which, through selection, the object can be said to ‘manifest itself *as such*’. In this way we can do justice to Horgan & Tienson’s claim that, when we perceive an object, ‘the experience is of a *temporal* object, an object that endures’ (2002: 521).

The operation of selecting an object so that one might go on to access each of its properties, in a serial fashion, is one that takes time. The success of this operation, then, depends upon the subject tracking one and the same object throughout. Otherwise the subject will end up accessing a different set of property instances (or tropes) to the ones they set out to access at the start of the operation. This is why, when we are aware of making no new selections, from the period t_1 through to t_5 , we are justified in believing that the object selected at t_5 is identical with the object selected at t_1 , even though many of the properties that we have visual access to may have changed during the intervening period (because, as might be the case, the object has altered during that period).

During any period in which *no* new selection is made, any changes in the subject’s direction of gaze, or in the properties that the subject has visual access to, will be due entirely either to changes intrinsic to the particular object, or to changes in the subject’s relation to that same object. Moreover, from her experience of the kinds of changes that a single selection can survive the subject may extrapolate to the identity conditions of objects. For example, she will learn that the only way her visual access can jump from one location to another, discontinuous location, without traversing the intervening locations (in such a way as to afford access to what is at those locations), is if she makes a new selection. On the other hand, she will also learn (or have verified), that an object may survive changes in location provided that these carve a continuous path through space. These two points are different sides of the same coin. Though more primitive creatures are likely to explicitly think (if they do at all) only of the second side of the coin.

Let’s suppose that the tiger that was focused on its prey as it padded down the veldt fails in its attempt to catch its prey. It is difficult to know what, if any, thoughts tigers

are capable of. But we can certainly imagine that the tiger might lament its missed opportunity after the prey has disappeared from sight. If the tiger does lament its missed opportunity then these thoughts clearly concern the very object that it successfully visually tracked throughout the failed attack. There is no reason whatsoever to say that the tiger's thoughts instead concern just some temporal stage of the prey²⁸ (or the collection of properties that define this temporal stage) - say, the temporal stage that finally evaded the tiger's lunge.

If the tiger is capable of learning from reflection on its mistakes then it may be that the best lesson to take from the episode is that the tiger ought to have approached from downwind, rather than from upwind. But if the tiger's thoughts concerned only the temporal stage of the prey that evaded the tiger's lunge then the tiger could not learn this lesson. Alternatively, rather than lamenting defeat the tiger might instead persevere with the hunt and think about where its prey is likely to go next (perhaps, if the river meanders and the prey is following the river, the tiger can cut its prey off). Again, if the tiger's thoughts concerned only temporal stages of the prey then no such planning would be possible.

These thoughts, which we have imagined the tiger to undergo, all exhibit sensitivity to the identity conditions of the relevant object – the prey. But it's worth noting that this does not require the tiger to think *about* identity conditions. Nor must the tiger think *about* patterns in her visual attention. The proposal is just that thoughts about particular objects are mediated (and justified) by an awareness of these patterns, but 'awareness of' does not mean 'thought about'. It's also worth noting that none of the thoughts considered above require that the tiger can think about objects *in general*. Perhaps there are other reasons for saying that one cannot think about a particular object unless one can think about objects in general, but none are to be derived directly from anything that I have said here.

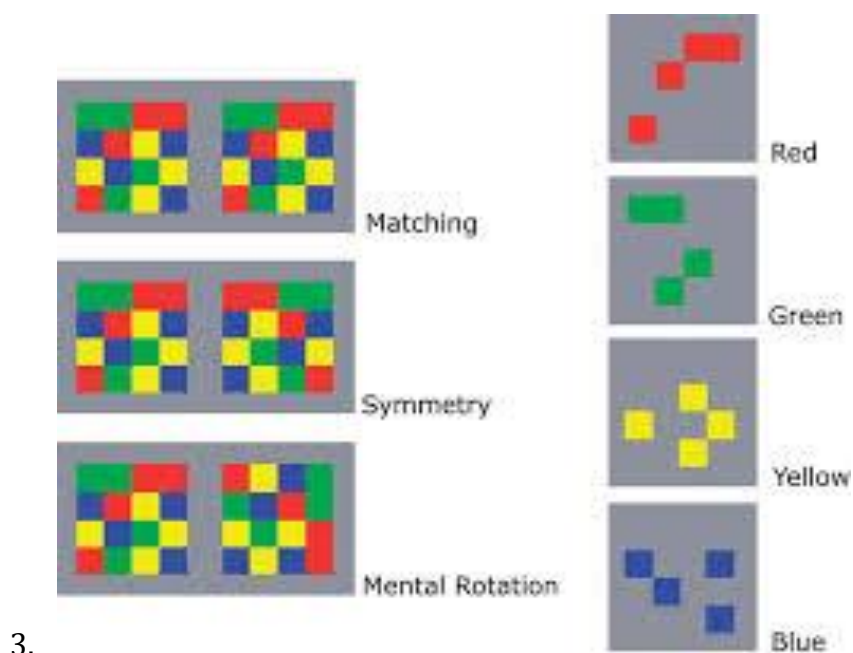
I have appealed to the agency that we exert over selection in arguing that it is highly plausible that we are generally aware of whenever a new selection has been made.

²⁸ Cf. Chapter 2 of Quine's *Word and Object* (2013).

But this fact is also the source of a potential objection to my proposal. We are capable of selecting regions of space that do not define any object at all. We are even capable of selecting discontinuous regions of space.

The diagram below illustrates three tasks. In the first task subjects must determine whether or not the images on left and right match one another. In the second task subjects must determine whether or not the image on the right is symmetrical with the image on the left. In the third task subjects must determine whether or not the image on the right is a rotation of the image on the left.

It has been established that, when performing any one of these three tasks, subjects must ascertain the answer to the relevant question, separately, for each of the colours in the display. For example, when subjects are asked, in the first task, whether the image on the right matches the image on the left, subjects must first assess whether the shape formed by all of the red squares on the left matches that of all of the red squares on the right. Then they repeat the process for the green squares, and so on. What subjects are not able to do is simply assess whether the image on the left matches the image on the right in one simple comparison.



3.

This diagram, printed by Huang and Pashler (2007), illustrates three tasks.

The regions defined by each colour are discontinuous, and they do not in any obvious sense define a single object. And yet we can (indeed, in order to perform the task, we must) visually select these regions. So this task makes vivid the fact that we are capable of selecting discontinuous regions of space, which do not seem to define a single object.

In addition to the bottom-up processes that I have already mentioned, which enable us to track selected objects as they move, there are also bottom-up processes that continually serve to orient us, and to orient the parameters that we use for selection, towards objects in the first place. The subject is able to override the dispositions and to instead select (perhaps spatially discontinuous) regions that do not define any single object. We do this when we select the regions occupied by each of the colours in figure x, in a serial fashion, in order to evaluate the overall figure for symmetry. But throughout the performance of such a task we must continually inhibit the disposition to instead select a 'genuine' object.

Even on the assumption that the purpose of selection is to present objects as such, we should not be surprised that we are capable of selecting spatially discontinuous regions of space, for objects are often partially occluded. Indeed, Moore, Yantis & Vaughan (1998) have demonstrated that attention *automatically* spreads from one region to another, spatially discontinuous, region, when these two regions appear (because of occlusion cues) to be parts of a single, partially occluded, object.

In the next chapter I shall consider the empirical evidence in favour of the two claims that I have made in support of the view that, through selection, the object manifests itself as such, and manifests itself as identical throughout the period during which it is selected. I will also turn to the question of how we should characterise those cases on which we do not select an (single) object (as in the matching/symmetry/rotation task, above).

Chapter 3 – Implications and Empirical Support.

In section 1 I outline some of the empirical literature supporting the view propounded in chapter 2. In section 2 I consider an implication of this view – that we only perceptually experience one object at a time -, and I deal with two potential objections that this might give rise to. In section 3 I offer an account of how we ought to characterise those cases in which we do not select an (single) object.

3.1: The Empirical Basis.

In section 2 of chapter 2 I made the following two claims:

1. Our visual systems orient us, and orient the parameters that we use for selection, towards objects.
2. Our visual systems dispose us to track those objects through space.

I now offer the empirical support for these claims. Yantis & Hillstrom (1994), and Hillstrom & Yantis (1994), have presented evidence that when a new object appears in the visual field this will capture our attention, at least briefly, even when this runs counter to our goals. Previous research has show that abrupt visual onsets capture attention. Hillstrom and Yantis (1994) wanted to establish whether the mechanism underlying this phenomenon was a) a luminance-change detection system, or b) a mechanism that detects the appearance of a new perceptual object. They distinguished these hypotheses by designing experiments in which a pre-existing object underwent a change in luminance (luminance-change only condition), and in which new objects appeared but were distinguished from their backgrounds only by texture, motion, or binocular disparity, and not by luminance (new object only condition). They established that ‘attention is captured in visual search by the appearance of a new perceptual object even when the object is equiluminant with its background and thus exhibits no luminance change when it appears’ (1994, 95). They

also found that ‘a highly salient luminance increment alone is not sufficient to capture attention’ (1994, 95).

In another experiment Yantis and Hillstrom (1994) established that motion can guide attention when motion is known by subjects to be ‘perfectly informative about the location of a visual search target, but that it does not draw attention when it does not predict the target’s position’ (1994, 399). They also found that ‘when motion segregated a local letter from its perceptual group, the local letter captured attention’ (1994, 399)²⁹ and they hypothesise that when motion serves to segregate a perceptual element from a perceptual group, a new perceptual object is created, and so attention will be drawn to it even when this runs counter to the subject’s goals.

Further evidence that, at a very early stage in visual processing, the visual system a) segregates the visual environment into discrete objects, and b) automatically tracks these objects as they move through space, comes from Kahneman, Treisman & Gibbs (1992) They made use of the ‘priming effect’, whereby the prior occurrence of a particular letter decreases the recognition time for that letter. Kahneman *et al.* showed that the priming effect for a letter traveled with the box in which it had originally occurred.

This offers indirect evidence in favour of both claims 1 and 2 above. The fact that, at a very early stage in visual processing, the visual system a) segregates the visual environment into discrete objects, means that the visual system at least has the right articulation of the visual environment in place in order for it to orient us, and orient the parameters that we use for selection, towards objects (i.e. claim 1 above). The fact that at a very early stage in visual processing, the visual system b) automatically tracks these objects as they move through space, shows that the visual system is capable of tracking objects through space in a bottom-up fashion, and so increases the plausibility that the visual system might likewise lock our attention onto an object once we have visually selected that object (i.e. claim 2 above).

²⁹ They call it a *local* letter because the perceptual groups, of which the local letters are members, are also letters (perhaps they would call the groups ‘*global* letters’).

There is a wealth of data now supporting some form of the idea that the unit of visual attention is the object (rather than a region of space). The first experiment that supported this idea came from Duncan (1984). He performed an experiment in which subjects were presented with a display consisting of a box with a line through it, and they had to make judgements either about two properties of the box, two properties of the line, or one property of each. Duncan found that 'two judgments that concern the same object can be made simultaneously without loss of accuracy, whereas two judgments that concern different objects cannot' (1984, 501). This supports the idea that, once we have selected an object, the visual system disposes us to stay locked onto that object. Moving our attention between objects, even when these roughly occupy the same region of space, comes with a cost.

More evidence comes from an experiment by Baylis & Driver (1993). Egly, Driver, & Rafal (1994) describe the experiment thus:

Baylis and Driver used an ambiguous display, analogous to Rubin's celebrated faces-vase figure, which could either be seen as two objects against a central background or as a central object against flanking background. Color instructions induced subjects to make either the two-object or single-object interpretation. Comparing the edges of the vase-faces was more difficult when these edges were seen as belonging to two objects (the faces) rather than one (the vase). This demonstrates a difficulty in simultaneously attending to two objects, and because an identical display was used in the one- and two-object conditions, the result cannot be explained in terms of differential locations or spatial frequencies (164).

Now, according to Huang and Pashler's (2007) Boolean map theory, in order to access properties instantiated by some part of an object, but not by the object as a whole, one must generate a Boolean map of the region occupied by just that part. Huang and Pashler's theory, on its own, does not predict that it would be more difficult to generate and compare Boolean maps when these pertain to what are seen as different objects than when they pertain to different parts of what is seen as a single object. But on the assumption that, once we have selected an object, the visual system serves to keep our attention locked onto that object, Baylis & Driver's (1993) findings make

sense. Perhaps selection of the object as a whole can survive brief interruptions, provided that these interruptions are due to the generation of a Boolean map for the purpose of accessing properties belonging to some constituent part of that same object.

The above suggestion - that selection of the object as a whole can survive brief interruptions, provided that these interruptions are due to the generation of a Boolean map pertaining to a part of the same object - means that we diverge from Huang & Pashler (2007). They say that selection is equivalent to the generation of a Boolean map. But it seems that we must say that the generation of a Boolean map pertaining to some part of an already selected object is not itself a case of selection.

Egly, Driver & Rafal (1994) performed an experiment in which subjects were presented with a display containing two rectangles. 'The task was to detect the "filling in" of one of the four ends of the two rectangles to yield a solid square at that end. Before the appearance of this square, one of the ends of a rectangle was brightened to induce covert orienting' (1994, 166). Now, as we know from Yantis and Hillstrom (1994), illumination alone is not sufficient to induce covert orienting. But in this experiment the illumination cue was strongly correlated with the subsequent location of the target, so subjects had good reason to orient towards the cue. This is consistent with Yantis and Hillstrom's (1994) finding that a luminance change will not 'capture' attention (where attention 'capture' is a purely automatic and bottom-up phenomenon).

Egly, Driver, and Rafal (1994) analysed the reaction times for 'invalid cue trials'. These are trials on which the location of the cue did not coincide with the subsequent location of the target. On these trials Egly et al. found that subjects were significantly slower in identifying the location of the target when it appeared in a different object (in the other rectangle) than when it appeared in the same object (at the other end of the same rectangle). This was despite the fact that in either case the target was the same distance from the invalid luminance cue.

This suggests that the cue caused subjects to select the rectangle that it appeared within, and that some mechanism in the visual system then disposed the subject to

stay locked onto that object. Above, in relation to Baylis & Driver's findings, I suggested that selection of the object as a whole might be able to survive brief interruptions, provided that these interruptions are due to the generation of a Boolean map pertaining to some part of the same object. If that's right then, on those occasions on which the target appears in the same rectangle as the prior invalid cue, subjects will be able to generate the relevant Boolean map – providing them with access to the location of the target - without having to select a new object. Egly et al.'s data provides further support for this assumption. And the interpretation of the data that this assumption makes available to us provides further support for the claim that, once we have selected an object, the visual system disposes us to stay locked onto that object.

That concludes my argument in favour of the claim that selection serves to present an object as identical throughout the period during which it is selected.

3.2 - Perceptual Experience of only one Object at a Time.

As Huang & Pashler put it, selection is binary; it distinguishes what is accessible from what is not accessible. And there cannot be multiple simultaneous selections, for, according to Huang & Pashler, the Boolean map – that is, one Boolean map – determines what the subject may access at any given moment. The whole point of the Boolean map is to describe those limitations on our attention that empirical investigation has discovered. If it were possible to make multiple simultaneous selections then it would be possible to transcend these limitations, and the Boolean map would not serve its purpose. The present proposal therefore entails that only one object can manifest itself in perceptual experience at a time.

That only one object can manifest itself in perceptual experience at a time might be thought problematic for a couple of reasons. One reason is introspective, and the other is empirical. The first issue is that, arguably, it introspectively seems as though we are simultaneously presented with an *array of objects*, stretching right across the visual field. But the force of this objection is diminished when we remember that the view being proposed *does* respect the intuition that we are aware of an array of

entities, stretching right across the visual field. It's just that the entities are only the *properties* of the objects that lie in the visual field. Since, at any moment, we're not disposed to track the objects to which many of those visible properties belong, we are not yet perceptually related to those objects. But as soon as we use some of those properties to select an object as figure from ground, we then become perceptually related to that object. This seems like a good middle ground, which takes account of the empirical evidence concerning the limitations of our perceptual contact with the world, but also respects our fundamental introspective intuitions.

The second issue is empirical. Pylyshyn & Storm (1988) have demonstrated that we have the capacity to track up to 5 objects in a display of 10 moving objects. This is a robust finding, and has been replicated many times since. Pylyshyn (1989) developed an account of this capacity that appeals to 'pre-attentive' representations of objects. But as Scholl, Pylyshyn & Feldman (2001) remark, the task certainly does *involve* attention. The implication seems to be that we can after all attend to multiple objects at the same time.

Yantis (1992) suggests that when we simultaneously track multiple objects, at the personal level we treat the elements as though they define a single, non-rigid object. On this interpretation of the task, at the personal level, we track a single virtual object through (sometimes dramatic) deformations in its shape. Of course, at the sub-personal level our visual systems must track multiple objects. It is only by doing this that we can use those tracked elements, at the personal level, to track the virtual object that they collectively define. But I am defending the claim that selection is the means by which objects manifest themselves in perceptual *experience*, so it is what goes on at the personal level that is relevant. Yantis performed a series of experiments in which:

“The ease with which the elements in the target set could be perceptually grouped was systematically manipulated. In Experiments 1-3, factors that influenced the initial formation of a perceptual group were manipulated; this affected performance, but only early in practice. In Experiments 4-7, factors that influenced the maintenance of a perceptual group during motion were

manipulated; this affected performance throughout practice' (1992, 295).

Factors that influenced the initial formation of a perceptual group included suggesting to the subjects, at the start of the experiment, that they might find it easier to track the elements if they imagine that the elements form a single object. Factors that influenced maintenance of a perceptual group pertained to the ways in which elements were constrained to move, relative to one another. Factors that influenced maintenance of a perceptual group affected participants' performance throughout the task. Factors that only influenced the initial formation of a perceptual group only affected performance in early practice, and this was taken to suggest that those manipulations that made the initial formation of a perceptual group harder simply delayed subjects' discovery of the strategy. Yantis' (1992) data suggests that our capacity to simultaneously track multiple objects does not undermine the claim that we only perceptually experience one object at a time.

3.3: Cases in which we do not Select an (Single) Object.

How must we characterise those occasions on which there is no single object that is selected? One tempting answer is to say that in such cases it is a region of space that is presented *as* an object. However, there is reason to resist this account of the situation, and to instead postulate that what is selected, and what is thereby presented *as* an object, is a 'virtual object'. The significance of the notion of a 'virtual object' will become clearer as we consider the reasons for positing it.

Sears & Pylyshyn (2000) performed an experiment designed to rule out the hypothesis that in multiple object-tracking (MOT) tasks subjects simply spread their attention across the smallest region that encompasses all of the target items. In a variant of the MOT paradigm Sears & Pylyshyn (2000) had subjects track 4 figure of eight items in a field of 8 randomly moving figure of eight items. At various points during the task one of the figure of eight items would transform into an 'E' figure or an 'H' figure, and subjects were to respond to such changes by indicating whether it was an 'E' figure or an 'H' figure. They found that subjects are quicker to respond to form changes in those items that they are tracking than they are to respond to form changes in 'distracter' items.

Importantly, the advantage appears to be target-specific: in particular, it doesn't apply even to those distractors that are located within the convex polygon bounded by the moving targets. If, when performing MOT tasks, subjects simply spread their attention across the smallest region that encompasses all of the target items, then one would expect quicker responses to form changes occurring to distractors when these occur within the convex polygon bounded by the moving targets.

‘In Intriligator's terms, these results all indicate that attention is split between the target objects rather than being spread among them’ (Scholl 2001, 10).

If the interpretation of this data, which Scholl attributes to Intriligator, were correct, then this would undermine the interpretation of Campbell's view that I have offered. For I have claimed that we can only attend to (and perceptually experience) one object at a time. But, according to Scholl, Intriligator interprets the above data as implying that ‘attention is split between the target objects’.

As I have already said, Pylyshyn considers the indexes implicated in MOT tasks to be ‘pre-attentive’. Sear's and Pylyshyn's (2000) interpretation of the data appeals to both pre-attentive indexes and to a ‘*unitary* focal attention’. They suggest that there is some kind of change detector that informs the subject of form changes occurring anywhere in the visual field, but that this detector does not inform the subject of the location of any such form change. This means that the subject must scan the visual field in order to locate the change and report whether the relevant item has changed to an ‘E’ figure or an ‘H’ figure. Their suggestion is that the indexing of items, which is required in order for the subject to track those items, provides a direct path for focal attention to travel along to those tracked items. They suggest that, for this reason, tracked items will always be checked first.

I earlier appealed to Yantis' (1992) data suggesting that when subject's perform MOT tasks, at the personal level, subjects treat the targets as though they define a single object, and they track this object through deformations in its shape. This view is consistent with Sear's and Pylyshyn's (2000) data, for Yantis' view is not that we select the region of *space* defined by the ‘virtual polygon’, but that we select a ‘*virtual*

object' defined by the 'virtual polygon' (and that this 'virtual object' just so happens to occupy that region of space).

It is well known that when two extended events, involving distinct sets of objects, are spatially superimposed over one another, subjects are capable of selectively visually attending to only one of these extended events (and to only the relevant set of objects). Neisser & Becklen (1975) were the first to establish this, and they also found that when subjects attend to just one of these extended events they tend to be oblivious to 'odd events' that occur in the other, unattended, extended event. This is despite the fact that both of the extended events occur in the same region of space. Given this capacity that we have to selectively attend to only one of two events occurring in the same region of space there is no reason, on Yantis' account, to expect that responses to form changes in distractor items would be any quicker just because these distractor items happened to occupy a location within the convex polygon bounded by the moving targets.

The tempting account of those cases in which we do not select a single object is that in such cases it is a region of space that is presented *as* an object. But the research from Sear's & Pylyshyn (2000), and from Yantis (1992), collectively suggests an alternative description of these cases: when we fail to select a genuine object, instead a 'virtual object' is selected, and it is this 'virtual object' that is thereby presented *as* an object. This virtual object is not something that can simply be reduced to the region of space that it occupies at any given moment, for there may be items that fall within that region but which the visual system does not treat as part of the virtual object.

We can treat these cases as analogous to imaginings of objects, or as analogous to hallucinations of objects. The selection is analogous to an imagining of an object if the subject intentionally inhibited those dispositions that orient us towards objects, and she knows that she has not selected an object (as in the MOT task, and the multi-coloured matching/symmetry/rotation task). The selection is analogous to a hallucination of an object if she did not intentionally inhibit the dispositions, but the manifestation of those dispositions simply failed to relate the subject to an object

(perhaps they related her to two distinct objects, which looked like a single object, or to a hologram), in which case she may not know that she has failed to select an object (at least, not immediately).

We have now established that Relationalism can explain our perception of objects, so this objection to the plausibility of the account is defused. From here on I will be mounting a positive argument in favour of Relationalism (or 'Naïve Realism', as I shall now return to calling it), over Intentionalism. This argument begins, in chapter 4, with an elucidation of some, as of yet unidentified, features of visual phenomenology. These are the original phenomenological observations, that will enable us to move beyond TE, and which were promised in the Introduction.

Chapter 4 - Perceptual Experience of Spatial Composition

In chapter 4 we shift our attention from visual experience of objects – the topic of chapters 2 and 3 – and onto visual experience of properties. I will identify an aspect of visual phenomenology, and a corresponding phenomenal content, which has until now gone unnoticed. The relevant content pertains to spatial composition relations. I will isolate the relevant aspect of perceptual experience by describing a number of capacities, which we do in fact possess, and which would be explained by perceptual experience of spatial composition. I will consider an alternative explanation of the capacities, and in responding to this I will lean on some claims that I make about the nature of the relevant phenomenology. I will also make the case that our phenomenology is consistent with the claim that we perceptually experience the *necessity* of spatial composition.

In what follows I shall talk of ‘phenomenal-shapes’, such as ‘phenomenal-squares’, and of ‘phenomenal-colours’, such as ‘phenomenal-red’. Wherever I talk of ‘phenomenal-’ and the hyphen is followed by a term that stands for a type of property or relation, I am talking about the type of element of phenomenal character – or the type of ‘phenomenal property’ - associated with perceptual experiences of the mind-independent property or relation that follows the hyphen. For example, phenomenal-red is the phenomenal property associated with perceptual experiences of the colour red.

My descriptions of the nature of the relevant phenomenology serve multiple purposes in this chapter. They offer direct support for the claim that spatial composition features in phenomenal content, they explain why this pervasive feature of phenomenology has so far gone unnoticed, and they support the additional claim that phenomenal-composition is a metaphysically necessary relation. This third claim is particularly important for the main line of argument that follows, in chapters 7-9. Given the central role that’s played by what is essentially a single phenomenological observation, throughout the chapter we shall return, over and over, to the

phenomenology of visual spatial experience. My hope is that the reader will become more acutely aware of the relevant, subtle, features of her visual phenomenology on each occasion that we return to them, as we successively approach the same introspective observations from slightly different conceptual perspectives.

So, my primary aim is to establish the following four claims:


- 1) Spatial composition features in phenomenal content and (consistent with Thesis 3), there is a corresponding element of phenomenal character, which I shall call 'phenomenal-composition'.
- 2) Phenomenal-composition is a metaphysically necessary relation. It arises from the intrinsic nature of the composed and the composing phenomenal-shapes. The fact that phenomenal-composition arises from the intrinsic nature of the composed and the composing phenomenal-shapes explains why it is so easy to miss (and why it so far *has* been missed). Though it is a distinct element of phenomenal character, it is not something *extra* in the sense that it could be imaginatively subtracted from the composed and the composing phenomenal-shapes.
- 3) Phenomenal-composition is not the product of cognitive penetration.
- 4) The *necessity* of phenomenal-composition is not the product of cognitive penetration.

In section 1 of this chapter I will make the initial case for the claim that spatial composition features in phenomenal content (claim one). In section 2 I will consider, and rebut, an objection to my argument in section one. In section 3) I explain how claim 2) is entailed by claim 1) together with a key observation that I used to motivate claim 1), and I offer some elucidation of and independent motivation for claim 2). I will also explain, in section 3), why phenomenal-composition cannot be the product of cognitive penetration (claim 3), and why the necessity of phenomenal-composition cannot be the product of cognitive penetration (claim 4). Finally, in

section 4), I consider a putative counter-example to claim 2 – Kanizsa shapes.

4.1. The Case for Spatial Composition in Phenomenal Content (Claim 1)

In this section I will present three arguments in favour of the claim that composition relations among spatial properties feature in phenomenal content. First, I will explain exactly what I am claiming.



4.

This image illustrates three objects of perception ³⁰.

I propose that whenever we perceptually experience a surface as being some shape – e.g. square –, we do so *by* perceptually experiencing the border of that shape. Moreover, we always perceptually experience the surface as square by perceptually experiencing the border of the shape – the shape instantiated by the surface – *as* composed of 4 straight lines. There are, then, 3 things that we visually experience in order to visually experience the surface as being square. These are: the border of the shape (that is, the shape instantiated by the relevant surface)³¹, the spatial properties

³⁰ Thank you, Rachel Dixon, for the above image.

³¹ Green (2017) offers a detailed, phenomenologically and empirically motivated, argument in favour of the claim that we perceptually experience objects as simultaneously instantiating multiple shapes. He would say that we perceptually experience the above (2-dimensional) object as simultaneously instantiating the ‘abstract’ property of being rectangular, and the ‘metric’ property of being square. The abstract property of being a rectangle is more ‘stable’ than the metric property of being square, because the property of being a rectangle can survive transformations that the property of being square cannot survive. In particular, Green argues that Marr’s 2 1/2 D sketch, which is just ‘an array

out of which the border of the shape is composed – in the case of the border of a square, this is 4 straight lines -, and the composition relation between these – e.g. the composition relation between the straight lines and the border of the square. This constitutes visual experience of the border *as* being composed of the straight lines. From here on, for ease of expression, I will talk of squares (rather than of their borders) as being composed of straight lines, and of them as being perceived as such. However, it must be remembered that this is used as short hand for the more precise formulation – involving borders –, which I have offered above.

The particular surface is perceived as square because the subject perceives the surface as instantiating squareness (and she does this in virtue of perceiving the surface as instantiating a property necessarily composed of 4 straight lines of equal length). The property, squareness, which the surface is perceived as instantiating, is an abstract object. In perceiving that property - that abstract object - *as* instantiated by a particular concrete object (see chapters 2 and 3 for an account of how particular concrete objects show up in perceptual experience, as such) one thereby perceives the property *instance* (or trope), *as such*³².

Still, one might feel some resistance to this claim, for it may seem that I am needlessly multiplying objects of experience. Really, one might object, once you've seen the location of the 4 straight lines, you get the square for free; the square isn't anything extra. And with this I am in agreement. But that is just to say that the square arises from the intrinsic nature of the located 4 straight lines – the relation between the located straight lines and the square is metaphysically necessary. Yet we may accept

specifying the viewer-centred distance, direction, and local orientation at each point (up to a certain resolution) for all visible surfaces in the scene' (2017, 359), cannot be a complete characterization of the content of perceptual experience. He also argues that Peacock's scenario content, an enrichment of Marr's basic idea, cannot be a complete characterization of the content of perceptual experience.

³² It might look as though this invites an objection that utilizes Bradley's Regress. However, recall that a property is perceived as instantiated by an object because selection of that object affords us with access to that property. So, there is no need to appeal to perception of an abstract object (i.e. the universal: the instantiation relation) in order to explain our experience of the property *as* instantiated by that object, and the regress never gets going.

this and still maintain that there is a distinction to be drawn between the 4 located straight lines, the arrangement that they collectively compose (the square), and the (metaphysically necessary) composition relation between the located straight lines and the arrangement. My claim is that we perceptually experience all three of these things.

I am saying that there exist phenomenal-squares, phenomenal-straight-lines, and phenomenal-composition. Phenomenal-composition is an easy feature of our phenomenology to miss, precisely because the relation between each of these elements of phenomenal character is so intimate. For I believe that just as the square is essentially related to the straight lines out of which it is composed, so to the phenomenal-square is essentially related to the phenomenal-straight-lines out of which it is phenomenally-composed. But this claim I will argue for in section 3 – for now I simply mention it in order to explain why the feature of phenomenology that I claim to exist is so easy to miss, and why it has never previously been identified. Of course, what I say about squares and straight lines applies to spatial part-whole relations more generally. In these more general terms, my claim is that we perceptually experience wholes, we perceptually experience their parts, and we perceptually experience the composition relation between these.

To further elucidate the sense in which composition relations manifest themselves in perceptual experience, I should point out that this is not simply a matter of perceiving the square to be co-located with the straight lines, for we also perceive shapes to be co-located with colours, but we are not tempted to think that there is any composition relation between shapes and colours. Again, we perceptually experience the relationship between squares and straight lines as tighter than that between a shape and some colour. That is, there is an additional kind of unity in our experience of spatial wholes and their parts, which we do not find in our experience of co-located spatial properties and colours.

It is in the nature of unities that they resist attempts at articulation, for articulation implies division. My hope is that the reader will get a clearer sense of what I am proposing by reading the arguments that I offer in its favour. For now, I simply urge that, despite the intimacy of these features of our phenomenology, still we must resist

the occasional temptation to conflate these experiences. We see the arrangement (the square), we see the located straight lines, and we see the arrangement *as* an arrangement of those very straight lines. So there are three distinct (though intimately related) things here that we perceptually experience.

4.1.1 Argument 1 - Perceptual Experience of Similarity Relations.

We are capable of perceptually experiencing similarities and differences between different types of spatial properties, and this capacity is best explained by the proposal that we perceptually experience spatial composition relations. For example, there is an important respect in which squares and triangles are similar to one another, and in which neither squares nor triangles are similar to circles: squares and triangles are both composed of straight-lines, and circles are not. So the explanation for the fact that we can perceptually experience a similarity between squares and triangles – one not shared by circles - is that we can perceptually experience the composition relation that holds between squares and triangles, on the one hand, and straight-lines, on the other.

Suppose that someone were to present you with a square, a triangle, and a circle, and to ask you which pair *looked* most similar. Suppose that you answer ‘the square and the triangle’. In response to which they ask: ‘why?’. You might well reply ‘they are both composed of straight-lines, unlike the circle’.

One might object as follows: we do not visually experience composition relations between colours, yet we can visually experience similarities and differences among the colours. If this is correct then we have reason to doubt that visual experience of similarity/difference relations between spatial properties requires perceptual experience of spatial composition relations. However, Allen (2016) has argued that we do in fact perceive certain colours – the ‘composite hues’ – as ‘phenomenal mixtures’, consisting of two or more of the ‘unique hues’. If this is correct then our capacity to perceptually experience similarities and differences among the colours is not a counter-example to the claim that perceptual experience of similarities and differences among property types (and so types of shapes) is best explained by perceptual experience of composition relations.

Nonetheless, the objector might maintain that Allen has a different notion of 'composition' in mind. For example, Allen (2016) says:

'Nor do the constituents of a phenomenal mixture exist as distinct, separable elements of the mixture: they are not like peas and beans in a stew (Hering 1920: 20). It is better to say that a phenomenally uncomposed colour like unique yellow appears neither reddish nor greenish nor bluish, whereas a phenomenally composed colour like orange appears both reddish and yellowish (Byrne and Hilbert 2008; cf. §6.3)' (2016: 142-143).

He also contrasts the composition of phenomenal-colours with that of the 'musical chords in which the constituent elements are still distinguishable (2016: 145). This demonstrates that Allen has a different kind of composition in mind, for in the case of spatial composition we do perceptually experience both the composing and the composed shape properties, and both are, to some extent, 'distinguishable'.

It does seem to me that the 'phenomenal mixture' involved in perceptual experience of colours – phenomenal-colour-composition, if you will - is importantly different from phenomenal-spatial-composition. Nonetheless, it seems to me that there is an abstract point of similarity between phenomenal-colour-composition and phenomenal-spatial-composition, and this abstract point of similarity is all that my argument requires. The abstract point of similarity is that:

S/D: Visual *experience* of similarities and differences among (mind-independent) property types manifests itself in similarities and differences in the intrinsic natures of the relevant phenomenal-properties.

I am happy to admit any account of perceptual experience that satisfies the above as an account on which we do perceptually experience composition relations. This doesn't rule out the possibility that there might be other ways of perceptually experiencing composition relations, which don't involve similarities and differences in the intrinsic natures of the relevant phenomenal-properties. But S/D does seem to accurately characterise *our* perceptual experience of composition relations – at least

when it comes to colours and spatial properties.

Anyone who thinks that we perceptually experience similarities and differences among property types and who is sympathetic to any kind of TE claim ought to endorse S/D. This is because if S/D were false then, assuming (as we are – see the discussion of Speaks in the introduction) that phenomenal properties are available to introspection, visual experience would seem to be characterised by two distinct sets of properties, only one of which characterised by the similarity relations in question. The difference in similarity relations would surely then provide a clear means of distinguishing between the two sets of properties, and visual experience would not (in any sense) *seem* transparent.

Following Shoemaker (2006) (who explicitly espouses TE), Chalmers (2006) says:

‘We attend to the phenomenal character of an experience by attending to the properties that objects in the world appear to have. An extension of this intuition suggests that we discern similarities and differences in phenomenal character by discerning similarities and differences in the properties that objects in the world seem to have’ (2006, 62).

Here, Chalmers is describing Shoemaker’s (2006) view. As Shoemaker himself puts it (in a way that assumes Intentionalism):

‘Similarity in the presenting manifests itself in represented similarity in what is represented, and in the absence of perceptual illusion requires that there be similarity in what is represented’ (2006, 475).

Though the emphasis in the above two passages is different from my own, these passages are entirely consonant with what I am saying. For if visual *experience* of similarities and differences among property types manifests itself in similarities and differences in the intrinsic natures of the relevant phenomenal-properties, then, given that we cannot discern (where this involves attention) in introspection two sets of properties (shapes *and* phenomenal-shapes, or colours *and* phenomenal-colours) it follows that we would discern similarities and differences in the one set of properties *by* discerning similarities and differences in the other set of properties. The difference

in emphasis is just that, in their analysis of our *attention* to the relevant sets of properties, Chalmers and Shoemaker are prioritising the mind-independent properties. But since only one set of properties shows up in introspection, it's difficult to see what could justify a priority in either direction.

So, perceptual experience of similarity relations between both colours and shapes manifests itself, phenomenologically, in the intrinsic natures of those phenomenal-colours and phenomenal-shapes. The perceived similarity between pink and red – not shared with green – manifests itself, phenomenologically, in the intrinsic natures of phenomenal-pink and the phenomenal-red. It does not seem that one could subtract the phenomenological similarity without changing the intrinsic nature of these phenomenal-colours. Likewise, the perceived similarity between a square and a triangle – not shared with a circle – manifests itself, phenomenologically, in the intrinsic natures of the phenomenal-square and the phenomenal-triangle. It does not seem that one could subtract the phenomenological similarity without changing the intrinsic nature of these phenomenal-shapes.

Speaking of perceptual experience of *colour* composition, Allen denies that 'the constituents of a phenomenal mixture exist as distinct, separable elements of the mixture'. The implication of the above point is that, even if it is the case that in perceptual experience of *spatial* composition 'the constituents of (the) phenomenal mixture (*do*) exist as distinct, separable elements of the mixture' (in a way in which they do not in the case of perception of colour composition), still the constituents of the phenomenal mixture must be *essentially* related to the mixture, in exactly the way that they are. In other words, the constituents – e.g. the phenomenal-straight-lines – are necessitated by the intrinsic nature of the mixture – e.g. the phenomenal-square.

In the case of colour experience, one experiences similarities and differences among the colours in virtue of similarities and differences among the phenomenal-colours. One perceptually experiences the similarity between an instance of red and an instance of pink in virtue of phenomenal-red being similar to phenomenal-pink (That is not to say that this will suffice for the subject to recognise/conceptualise that the instance of pink is similar to the instance of red. This would additionally require that

the subject attend to the colours and to their relations to one another, and that the subject possess the relevant concepts).

The above suggests that it may be possible to run my argument with reference to perceptual experience of colour composition, as well as to spatial composition. The reason that I have chosen to run the argument with shapes is two-fold. The first reason is that it is possible to deny that similarity relations among phenomenal-colours reflect mind-independent similarity relations³³. Indeed, some people take this as reason for denying that colours themselves are mind-independent properties (for example, there are those who hold that colours are dispositions to cause certain types of experiences in observers)³⁴. The second reason is that there are two additional arguments for the claim that we perceptually experience spatial composition relations, at least one of which (argument 3) does not appear to apply to the colours.

4.1.2 Argument 2 – Perceptual Experience of the *Ways* in Which Shapes are Similar

The second argument for the claim that composition relations among spatial properties are perceptually experienced is that if spatial composition were not something that we could visually experience then one might expect us to be unable to provide any precise answer to the question ‘in *what way* do the square and the triangle look more similar to one another than either do to the circle?’ Yet it is quite plausible that someone would very quickly answer such a question with ‘they are

³³ See Shoemaker (2003), who denies that the similarity relations among those elements of phenomenal character associated with colour experience constitute perceptual experience of any mind-independent similarity relations, but who does not hold an error theory (this is because he then held that they constituted perceptual experience of similarity relations among ‘appearance properties’, rather than between colours, and that perceptual experience is not misleading in this regard).

³⁴ See Johnston (1992) for an endorsement of this view for reasons along these lines. Johnston does not argue from the claim that the similarity relations among the colours actually do not reflect the similarity relations among any mind-independent properties. Johnston argues from the claim that we can tell on the basis of perceptual experience alone that the colours do instantiate certain similarity relations, but we cannot tell on the basis of perceptual experience alone whether or not any mind-independent properties instantiate those similarity relations. But it is possible to concur with Johnston’s first claim, and yet resist his second claim. Allen (2016) does precisely this.

both composed of straight-lines, unlike the circle'. Visual experience of spatial composition explains the fact that we are able to say, with such speed and precision, *how* shapes look similar and different to one another.

It is true that squares and triangles just are ways in which straight-lines can be arranged, but if this fact is to explain how we are able to provide such quick and precise answers to questions about which spatial properties look most similar to one another then it must be the case that we visually experience the square (and the triangle) *as* an arrangement *of* straight-lines. That is, it must be the case that we visually experience the relation between the arrangement (the square) and the spatial properties (straight-lines) thus arranged. This is just another way of saying that we must visually experience the composition of the square (and the triangle), by the straight-lines. So, relations of spatial composition must enter into the content of visual experience.

If one were presented with three colour samples – blue, red, and orange -, and one were asked which colours *looked* most similar, one would immediately respond 'the red and the orange'. However, if one were asked why, then a precise answer might well prove more elusive. One might have to resort to metaphor: 'the red and the orange are warm, but the blue is cold', for example. Alternatively, if one did rapidly respond with something along the lines of 'the orange has some red in it, but it has no blue in it', this might well be evidence that the perceived composition implicated by the 'phenomenal mixture' that Allen (2016) speaks of, is, after all, more similar (in relevant respects) to the type of composition that obtains between spatial properties. In which case we have a further reason to believe that colour is not a counter-example to the claim that perceptual experience of similarity/difference relations among (spatial) properties requires perceptual experience of composition relations among (spatial) properties.

Just the fact that we are capable of perceptually experiencing determinable properties, as well as determinate properties, evidences what I am talking about. To perceive two different types of shapes – a short rectangle and a long rectangle, or an isosceles triangle and an equilateral triangle - as being instances of the same determinable – rectangle, triangle - is to perceive what are actually two different

types of properties as being in some way similar. We perceptually experience short rectangles and long rectangles as both being rectangles in part because we perceptually experience both as composed of straight lines. It is crucial here that we can perceptually experience not only that the two different rectangles are similar in some way, but also *how* they are similar, for this determines how we group different determinate properties into determinables.

The above phenomenon is exactly the same phenomenon as that which I previously described: when we perceive a square as being similar to a triangle. The only difference between these cases is that when it comes to those properties that are more commonly classified as determinables, we have a predicate that applies to all those shapes that are similar in the right way. But we could have a predicate for all those shapes that are composed of straight lines, and if we did, then we would say that triangles and squares – but not circles - are instances of that same determinable property³⁵.

One might be tempted to claim that it only seems as though we must perceptually experience composition relations because I have neglected the fact that perceptual experience is typically of objects. A single object can at once be represented as being both a rectangle and a short-rectangle. In virtue of representing it as a rectangle, I am able to perceptually experience it as similar, in one respect, to all other rectangular objects. In virtue of representing it as a short-rectangle, I am able to perceptually experience it as different, in one respect, from all other objects not represented as short-rectangles (including some objects which are represented as rectangles). This is all we need.

The problem with the above suggestion is that it patently is the case that we can, in

³⁵ It's worth noting, though, that one needn't *attend* to the composition relations in order to *attend* to the determinable property. As I will argue in the next chapter, it seems that our concepts of spatial properties can be observational concepts, where possession of such concepts requires only the capacity to recognise instances. In which case, it is possible to possess the concept of a square without possessing the concept of straight lines or of composition relations, and without having any knowledge regarding the spatial properties that squares are composed of. This would seem unlikely if attending to determinable properties required attending to the relevant composition relations.

fact, perceptually experience a determinate shape property, a, belonging to an object, b, as being in some way similar to a type-different determinate shape property, c, belonging to another object, d. We do perceptually experience similarities and differences between *properties*, including shape properties, and not just between objects. And, as has been mentioned, we are able to provide quick and precise answers to questions of the form ‘*why* do these two shapes (not objects) look more similar to one another than either do to this third shape?’ If we perceptually experienced objects as having myriad different shape properties of different determinacies, but perceptually experienced no composition relations between properties, then it seems likely that we would be unable to do this.

Green argues for a layered view of shape perception:

‘Visual representations of geometrical properties are layered in a hierarchy roughly in accordance with the stability of those properties. Thus, when you see a triangular surface of an object, your visual system constructs numerous representations arranged in a multi-level hierarchy: the object is represented at one level as having a quite specific metric shape (for example, a surface composed of points such-and-such a distance away with such-and- such orientation relative to the line of sight), but it is also represented at another level as a triangle, and at a third level as a solid (filled) figure’ (Green, 2017: 378).

There is an ambiguity in the passage above. When Green says that the object is perceived (Green assumes that this is a matter of representation) as a triangle, does he mean that it is represented as a triangle simpliciter, or as some more determinate type of triangle? Green certainly does think that we perceive the object as a triangle simpliciter, for he offers triangularity as an example of an ‘abstract’ property that we perceive – where an abstract property is a property that is capable of surviving certain kinds of transformations³⁶ -, and triangularity is abstract only if what is meant

³⁶ More specifically, Green considers shapes to be ‘abstract’ if they can survive affine, or even topological transformations. Though the distinction between ‘metric’ and ‘abstract’ properties is a relative one. For example, an *isosceles* triangle, which cannot survive even affine transformations, can survive translation and scaling.

is triangularity *simpliciter*. But I suggest that we need not, and ought not to, choose between these options.

We perceive the object as a triangle *simpliciter*, and as a triangle of some more determinate type. We can explain both of these perceptual accomplishments by drawing on composition relations between spatial properties, and by insisting that these composition relations are also perceptually experienced. We perceive the triangle *as* a triangle *simpliciter* because we perceive it *as* a property composed of three straight lines. And since we also perceive the angles that these straight lines subserve, we can say that we perceive the triangle *as* a triangle of some more determinate type because we perceive it *as* a property composed of straight lines that subserve the angles that they do.

4.1.3 Argument 3 – Our Capacity to Purposefully Manipulate Shapes.

The third argument for the claim that composition relations among spatial properties are perceptually experienced is that it seems that we can *see* how to rearrange the parts of an object's shape so as to influence its shape in a certain way. Arguably, if one has command of one's hands and one knows what a cuboid is, one is in a position to mould a ball of clay into a rough cuboid. In contrast, one does not know, simply by looking, how to alter the surface of an object so as to produce a particular colour, even roughly. Few of us really know this, because it requires knowledge of, for example, certain chemical theories and the functioning of the human visual system. Most of us do know how to mix paints so as to produce a particular colour, roughly. But this is a heuristic that we learnt (usually) in childhood, and long after we were able to purposefully alter shapes.


I remember having to consult my memory about which colours would combine to produce a specified target colour, but I do not remember ever having had to consult my memory about which shapes to combine, in which ways, in order to produce a specified target shape. One can imagine having to consult one's memory if the shape requested was a very complex one, but in this case it seems that what one is attempting to remember is simply which shape is in question. If one has learnt to visually recognise a very complex shape, but one is still not very practiced at

producing the shape, then one will likely have to consult one's memory if one is to produce that shape. But it seems that one will not be recalling *rules* for producing the shape from other shapes, rather, as I've said, one will simply be recalling which shape is in question. This is made vivid by the fact that if one was presented with an example of the shape, one would not need to consult one's memory at all, one would simply: look and copy. In contrast, if a child of a certain age were presented with an example of the target colour, she would certainly still need to consult her memory if she were then to produce it from other colours.

It is clear that children are adept at replicating target shapes, for there is a children's game that utilises this very capacity. Below are two images displaying shapes that can be composed with Tangrams³⁷.

'The **tangram** (Chinese: 七巧板; pinyin: *qīqiǎobǎn*; literally: 'seven boards of skill') is a dissection puzzle consisting of seven flat shapes, called *tans*, which are put together to form shapes. The objective of the puzzle is to form a specific shape (given only an outline or silhouette) using all seven pieces, which can not overlap. It is reputed to have been invented in China during the Song Dynasty,^[1] and then carried over to Europe by trading ships in the early 19th century' (Wikipedia).

³⁷ Thank you, Louise Richardson, for drawing my attention to this game.



5.

shutterstock.com • 1277518819

Two images displaying shapes that can be composed with Tangrams.

It seems that when one looks at a shape with an aim to replicate it, one really does *see* how the shape is composed. As one manipulates and orients 4 straight lines with one's hands, one *sees* the composition occurring before one's very eyes. Once the square has been composed, one *sees* how each of the straight lines contributes to the square, and one *sees* how the straight lines collectively compose the square. This is why it is so easy to replicate a seen shape (in contrast to the replication of seen colours).

4.2 An Alternative Explanation

I have argued, in section 1, that we perceptually experience spatial composition, on the basis of certain capacities that we have for manipulating shapes, and for making certain kinds of similarity judgements. But there is an alternative explanation of these capacities available. One might propose that:

Seeing-*that* a square bears composition relations to straight-lines would do just as well as actually *seeing* the composition relations.

To be clear, the suggestion is *not* that perceptual experience of spatial composition might be the product of cognitive penetration³⁸, for this would still imply that we perceptually experience spatial composition, and that is precisely the claim that I am arguing in favour of. In section 4 I will explain why perceptual experience of spatial composition cannot be the product of cognitive penetration. But the suggestion that I am arguing against at this point is simply that there is not perceptual experience of spatial composition *at all*. On this view, we merely perceive-*that* the shapes we see compose one another. That is, we know that the shapes that we perceive are composed out of each other on the basis of our perceptual experience of those shapes. But this knowledge does not have any influence on the phenomenal character of the perceptual experience.

The disagreement between someone who believes that we perceptually experience spatial composition and someone who believes that we merely perceive-*that* one shape is composed of others hinges on the following claim:

P) There is phenomenal-composition.

³⁸ See Siegel (2016) for a discussion of cognitive penetration and related phenomena.

Those who believe that we perceptually experience spatial composition will endorse P), while those who believe that we merely perceive-that one shape is composed of others will deny P).

4.2.1 An Objection – Merely Perceiving-*That* the Square is Composed of the Straight Lines.

The idea here is that we know that the shapes that we perceive are composed out of each other because we already have certain beliefs about which types of composition relations obtain between certain types of spatial properties when these properties are co-located. When we see those types of spatial properties co-located then we immediately arrive at the belief that they bear the relevant composition relations to one another. This explains our capacity to make the similarity judgements, and to manipulate shapes, and there is no need to say that we perceptually experience the composition relations.

One prima facie reason for rejecting P) above, and for endorsing the seeing-*that* only view, is that until now no one has ever identified P). The seeing-*that* only view offers to explain our capacity to make the relevant similarity judgements and to manipulate shapes, without proposing that there is some utterly pervasive feature of our visual phenomenology that has somehow been missed until now.

In response, I think that there are very good reasons why these features have evaded detection until now. It took us some time, at the beginning of this chapter, to clearly distinguish the properties that I am claiming that we are perceptually sensitive to. This is because these mind-independent properties are intimately related to one another, and this intimacy is reflected in our experience of them. The phenomenal properties that manifest the relevant mind-independent properties in our perceptual experience are, like those mind-independent properties, intimately related to one another. And as I said earlier, unities resist articulation, for articulation implies division. I will say a little more to elucidate this.

When we visually experience a square object we simultaneously perceptually experience several properties: a square, a colour, and 4 straight lines, all of which occupying roughly the same area. There is a kind of unity in our experience of the square, the straight lines, and the colour, afforded by the fact that they are all perceived as occupying roughly the same location. But there is a difference between the unity of our experience of the colour and the other properties, and the unity of our experience of just the straight lines and the square. There is an additional kind of unity in our experience of the straight lines and the square, which is not present in our experience of the square and our experience of the colour.

The additional kind of unity in our experience of the straight lines and the square, which is not present in our experience of the square and our experience of the colour, is brought out by the fact that one can visualise the phenomenal-square with a different phenomenal-colour, but one cannot visualise the phenomenal-square without phenomenal-straight-lines. What makes this fact pertinent is that this fact - about what it is possible to visualise - *seems* to be rooted in the nature of the relevant elements of phenomenal character. I will elucidate, and motivate, this claim by way of a contrast.

There are some things that it is impossible to visualise, but where there is no temptation to say that this impossibility has anything to do with the intrinsic nature of the relevant elements of phenomenal character. For example, plausibly it is impossible for a human to simultaneously visualise a football stadium full of thousands of faces and to visualise every single visible detail of every single face. Of course, I could be wrong about this. But the fact that there is an element of doubt over this claim serves to make my point. It seems to be a claim about the limitations of human visualisation, rather than a claim about what experiences are and are not metaphysically possible.

A contrasting example comes in the form of the Penrose Triangle, below.



6.

The Penrose Triangle.

In order to see this triangle as a three-dimensional shape one must trace one's eyes along the arms of the triangle. It's possible to see each vertex, individually, as exemplifying a rigid three-dimensional shape. But when one's eyes move from one vertex to another, there is a break with one's expectations, and it does not seem to be possible to synthesise one's visual experiences of all of the vertices. Moreover, even as one focuses on a single vertex there is an awareness of something in one's peripheral vision that somehow does not meet one's expectations. And one's awareness of this discordant feature of the scene seems to interfere with the task of even seeing the individual vertex as exemplifying a rigid three-dimensional shape. It is only by ignoring what is in one's peripheral vision that one is able to see the vertex as exemplifying a rigid three-dimensional shape. So it does not seem to be possible to synthesise these experiences into a single experience of a three-dimensional object, and, moreover, in order to even see a part of the object as exemplifying a rigid three-dimensional shape one must actively ignore the rest of the object.

Likewise, one can *imagine* going through the same visual routine, and one can thereby successively visualise each of the vertices. Possibly, with a little difficulty, one can even imagine one's awareness of the discordant features in one's peripheral vision. Certainly one can imagine the struggle to ignore one's peripheral vision. But what one cannot do is visualise the Penrose Triangle, as a three-dimensional whole

(though one can visualise it as a two-dimensional triangular whole, just as one is able to see it as a two-dimensional triangular whole). This seems to be precluded by the very nature of the relevant elements of phenomenal character (whether this be because the elements of phenomenal character are actually instantiated by the visualising, or they are merely imagined within it³⁹).

Similarly, when one claims that it is impossible to visualise a square (or triangle) without straight lines, this seems to be a claim about what experiences are possible. The claim about a square without straight lines, and about the Penrose Triangle, is rooted in an appreciation of the natures of the relevant elements of phenomenal character, in a way in which the corresponding claim about visualising a football stadium full of faces is not. Reflection on the natures of the elements of phenomenal character involved seems to make it evident that experience as of a square without straight lines, or as of a Penrose Triangle, are *metaphysically* impossible.

In both cases the problem is the same: a given type of phenomenal-whole necessitates specific types of phenomenal-parts. In the case of the Penrose Triangle, the phenomenal-parts are not of the right types, so there is no phenomenal-whole (that is, none of the right kind – one that corresponds to a rigid three-dimensional shape). In the case of a phenomenal-square without phenomenal-straight-lines, the phenomenal-parts are entirely absent, so there is no phenomenal-whole.

I claim that the composition relation manifests itself through this phenomenal unity of the phenomenal-square and the phenomenal-straight-lines. In this way visual experience is not only sensitive to the existence of the composition relation, but it is also sensitive to the profound intimacy of that relation. But in order that we might be perceptually sensitive to the intimacy of that relation it is crucial that the phenomenal relation between phenomenal squares and phenomenal-straight-lines be correspondingly intimate – this phenomenal unity/intimacy just *is* the perceptual manifestation of the corresponding mind-independent unity/intimacy. Perversely, then, it is actually the visual system's pervasive sensitivity to the intrinsic nature of

³⁹ See Martin (2002) for an argument that visualizing involves only the latter.

our environment that tempts us to altogether elide phenomenal-squares and the phenomenal-straight-lines from which they are phenomenally-composed, and to underestimate the features to which our visual experience is sensitive.

These considerations remove the putative advantage of the seeing-*that* only view. There is nothing mysterious about the fact that P) has been overlooked for so long, so there is no *need* for an alternative explanation of the relevant capacities which avoids commitment to P). The above considerations, regarding unity phenomenology, also make important steps towards providing a positive motivation for P). But in order to see why the seeing-*that* only view is so implausible, given phenomenological observations of the kind described above, we must consider the seeing-*that* only view in more detail. This we shall do in the following sub-section.

4.2.2 Perceptual Uses and Evidential Uses of the Language of Appearing.

The advocate of the seeing-*that* hypothesis must say that, although we are inclined to say things like ‘the square and the triangle appear more similar to one another than either do to the circle’, this use of the word ‘appear’ is to be explained away. AD. Smith distinguishes between: genuinely perceptual uses of the language of appearing, and merely evidential uses of the language of appearing (page 35 on).

‘To say that the economy seems to be improving, or that the military situation looks bad, is not to give voice to the nature of one’s perceptual experience at all. I call such uses merely evidential because they simply have the force of “the evidence points towards the conclusion that...” Such uses are purely judgemental, and typically express a belief that falls short of certainty.’ (Smith: 2002, 37).

Though A.D. Smith acknowledges that there are cases that might initially strike us as perceptual uses of the language of appearing, but which are in fact still evidential. He calls these “perceptual-evidential”. He offers the following example “the island looks inhabited’, said by the shipwrecked sailor who has noticed smoke in the distance and

footprint-like marks in the sand' (2002, 38). He contrasts such perceptual-evidential uses of 'appears' with genuinely perceptual uses, such as 'that wall appears yellow'.

What is distinctive of evidential uses of 'appears' is that we are in a position to offer independent evidence for the judgement. This means evidence for the judgement 'o looks F' other than its simply looking F. For example, the sailor can appeal to the smoke and the marks in the sand as evidence for the judgement that the island is inhabited. What is distinctive of genuinely perceptual uses of 'appears' is that 'the focus of our judgement is the very evidence of our senses... that the wall looks specifically yellow to you, rather than some other colour, is not something for which you have any independent evidence at all. It just does look yellow' (2002, 38-39).

The proponent of the seeing-*that* hypothesis claims that the case in which the subject says 'the square appears more similar to the triangle than either do to the circle' is a merely evidential, or perceptual-evidential, use of the word 'appears'. But this doesn't sound right. What is our independent evidence for the judgement that the square is in some respect more similar to the triangle than either is to the circle? Surely we have no independent evidence. It just does look more similar!

One might object that I have picked examples of evidential uses of 'appears' that are particularly unlike the case in hand, in order to dissuade the reader from assimilating 'appearance' talk regarding composition, to evidential uses. Perhaps a better example – one that helps to highlight the affinity of 'appearance' talk regarding composition to evidential uses – is the case in which someone *appears* happy. Clearly we do frequently perceive-*that* people are happy, but few would want to say that we literally perceive their happiness⁴⁰. When we say that someone appears happy we are making an evidential use of the word 'appears'. But even in this case it is possible to discern independent evidence for the judgement. As may be the case, the evidence is that the person is smiling and we believe that smiling is indicative of happiness.

⁴⁰ McDowell (1982) is sympathetic to the view that we do literally perceive others' 'inner' states. Though he seems, in the end, more sympathetic to the weaker claim that we perceive 'his giving expression to his being in that 'inner' state' (1982, 387).

There does not seem to be any equivalent independent evidence for the judgement 'the square appears more similar to the triangle than either do to the circle'.

As Smith puts it:

'The central point here is that for some, but only some, values of F, something can appear exactly like something that is F when it is not, *even though we are perceiving the thing veridically.*' (49)

Where the 'seeming', or 'appearing', is merely evidential, one can consistently believe both that one's perceptual experience was veridical, and that things were not as they seemed. The person who appeared happy may not have been happy, and yet one's perceptual experience of them may have been entirely veridical; they may have been putting on a brave face! The question at issue is whether the 'seems' in 'the square and the triangle *seem* more similar to one another than either do to the circle', is such that the square could fail to be composed of straight lines, and yet the perceptual experience is entirely veridical. This does not seem to be the case.

The proponent of the seeing-*that* only view has no option but to say the following. Given that we are in possession of the *general* knowledge that whenever 4 straight lines are co-located with a square those straight lines compose the square, on those occasions on which we do perceive 4 straight lines co-located with a square, we are therefore justified in judging that the particular square is composed of the particular straight lines. It is against the background of this general piece of knowledge that the perceptual experience of the 4 straight lines as co-located with the square constitutes *evidence* for the judgement that the square is composed of the 4 straight lines. So, when we say that the square 'appears' to be composed of the straight lines, what we mean is that the evidence points towards this fact.

This line of thought can be extended to account for judgements concerning the relevant similarity relations. First we judge that the square and the triangle are composed of the co-located straight lines, and that the circle is not composed of any straight lines, and then one judges that in virtue of this the square and the triangle are

more similar to one another than either are to the circle. Note, however, that on the seeing-*that* only view the subject judges that the square and the triangle are more similar to one another than either are to the circle. Can the subject also judge that the square and the triangle 'look' or 'appear' more similar to one another than either do to the circle? Yes, but only if the words 'look' and 'appear' are taken in the evidential sense. When the subject says that the square and the triangle look more similar to one another than either look to the circle, what she means is that the evidence points towards the square and the triangle *being* more similar to one another, in some respect, than either are to the circle.

On this view the claim that the square and the triangle *look* more similar to one another than either do to the circle is a claim about the perceived shapes, but it is not at the same time a claim about their phenomenal *looks* – no commitment is made to the phenomenal *looks* of the square and the triangle being in any way similar to one another. The claim is consistent, for example, with it being false that phenomenal-squares are in any way similar to phenomenal-triangles.

One option for the advocate of the seeing-*that* only view is to deny that there are similarity relations among the phenomenal-shapes that reflect any similarity relations among the shapes simpliciter. This advocate of the view will deny that there is such a thing as phenomenal-composition. They will reject the descriptions of unity phenomenology that I offered in the previous sub-section. And they will resist the explanation, offered in the previous sub-section, of our intuitions about the impossibility of visualising a square without visualising straight lines. Perhaps they will say that we are wrong to think that the case is any different from the case of visualising a football stadium full of faces, replete with details – in both cases we simply find ourselves unable to perform the task. The trouble is that this view is flatly at odds with introspection, and it runs roughshod over what seem to be genuine distinctions.

Alternatively, the advocate of the seeing-*that* only view may concede that there are similarity relations among the phenomenal-shapes that reflect similarity relations among the shapes simpliciter, but deny that by virtue of these we perceptually

experience the similarity relations among the shapes. They will accept that what I have been calling phenomenal-composition exists, but they will claim that phenomenal-composition, along with the similarity relations among phenomenal-shapes which phenomenal-composition underpins, serve no epistemic purpose. On this view, phenomenal-composition is epistemically redundant.

The problem with this view is that, given that these composition/similarity relations among phenomenal-shapes do mirror the composition/similarity relations that hold between the corresponding mind-independent shapes, what is to stop the subject from making use of them? After all, these composition/similarity relations are introspectively accessible to the subject, so she *could* use them as a guide when making her similarity judgements, and when manipulating the shapes of objects. And once one accepts that the subject can use the similarity relations among the phenomenal-shapes in this way then it just seems dogmatic to go on denying that they constitute perceptual experience of the similarity relations among mind-independent shapes.

It might seem that a way out of this predicament for the advocate of the seeing-*that* only view is to just slightly weaken the view, and to allow that when one sees-*that* the square is composed of straight lines this influences the phenomenal character of the perceptual experience, so that you also come to perceptually experience the composition relation. This is what writers have called 'cognitive penetration'. This concession would preserve some primacy for seeing-*that*. But my claim, at this point, is only that we do perceptually experience spatial composition relations, so this modification of the seeing-*that* view concedes exactly as much as I require⁴¹.

I think that the case against the seeing-*that* only view is already sufficiently strong. In chapter 5 I will seek to strengthen it a little further by examining more closely the commitments of the view. If the seeing-*that* (only) view were correct then our knowledge of composition relations in general would not have its source in

⁴¹ Though I will explain why phenomenal-composition cannot be the product of cognitive penetration later in the chapter.

perceptual experience, for on this view one can only perceive-*that* a particular square is composed of some particular set of straight lines because one already knows that squares in general are composed of straight lines. It is only against the backdrop of the knowledge that squares are always composed of those straight lines with which they are co-located that the perceived co-location provides evidence for the judgement of a particular composition relation. This, however, raises the question of from where this background knowledge originates.

In chapter 5 we shall consider the two options available to the proponent of the seeing-*that* only view. As shall become clear, neither of them is attractive. The first – that statements about spatial composition relations are analytic truths – comes under immense pressure from two different directions. The second option – that beliefs about spatial composition relations are innate – presents some intuitive and theoretical issues, and also forces us into a fresh confrontation with the phenomenological observations already made in this chapter.

For now, I will not pause any longer over claim 1. Instead, in the following section, I turn to claim 2 – that phenomenal-composition is a (metaphysically) necessary relation. However, there is a point of contact between what I will say in the next section and some of the arguments that I will present in chapter 6. Some of the accounts available to the advocate of the seeing-*that* only view of the source of our beliefs regarding composition relations fail to account for the particular way in which we think about the *necessity* of composition relations. More specifically, they fail to explain why we think of spatial composition relations as not merely *empirically* necessary, but as *metaphysically* necessary. In the following section I will argue that not only is it *phenomenologically* plausible that phenomenal-composition is a metaphysically necessary relation, but that once we acknowledge this then we can avail ourselves of a straightforward explanation of why it is that we think of spatial composition relations as *metaphysically* necessary relations.

4.3. The Necessity of Phenomenal-Composition (Claim Two), and Ruling out Cognitive Penetration (Claim 3).

In chapter 6 I will argue that, not only do we perceptually experience spatial composition relations, we also perceptually experience these relations *as* necessary relations. Once this has been established we can then, in chapters 7-9, use perceptual experience of the necessity of composition relations as a measure of the viability of competing accounts of perceptual experience. To the extent that a view of perceptual experience can accommodate perceptual experience of the necessity of spatial composition relations, in a way that is consistent with the phenomenology, that will count in favour of the view. Conversely, to the extent that a view of perceptual experience has trouble accommodating perceptual experience of the necessity of spatial composition relations, in a way that is consistent with the phenomenology, that will count against that view.

However, were it not for claim 2, none of the views of content fixation would have any trouble accommodating perceptual experience of the necessity of spatial composition. It is the requirement that views of perceptual experience accommodate perceptual experience of spatial composition *while also* adhering to claim 2 that favours certain views over others. In this section I shall defend claim 2, and in doing so we shall also see why phenomenal-composition cannot be the product of cognitive penetration (Claim 3).

Claim 2 is:

2) Phenomenal-composition is a metaphysically necessary relation. It arises from the intrinsic nature of the composed and the composing phenomenal-shapes. As phenomenal-composition arises from the intrinsic nature of the composed and the composing phenomenal-shapes, it's easy to miss – for

instance, it's not something *extra*, which could be subtracted from the composed and the composing phenomenal-shapes⁴².

As I shall explain, some of what I have already said in support of claim 1 also supports claim 2. In all, postulation of the tight connection between phenomenal-wholes and phenomenal-parts, described in claim 2, has three points in its favour. The first two of these, which I have already argued in favour of, are:

- a) Once one has a clear view of the distinction between a phenomenal-square, and the phenomenal-straight-lines out of which it is phenomenally-composed, introspection seems to support 2).
- b) Claim 2 helps us to understand how it is that phenomenal-composition has been overlooked until now.

I have already offered reasons for endorsing a) and b) in sections 1.1 and 2.1. One important point advanced in section 2.1, in favour of claim a), was that we seem to have different attitudes about the source and nature of our difficulties in visualising certain types of phenomena. More specifically, it seems intuitive that the task of visualising a square without visualising any straight lines presents a difficulty of a different kind to the task of visualising a football stadium full of faces, replete with facial details, all at once. While both tasks appear to be impossible, it seems that the first task is impossible in a more absolute sense. It seems that there is something about the nature of phenomenal-squares and phenomenal-straight-lines that precludes our fulfilment of the second task (and would preclude its fulfilment for any creature whose perceptual experience has phenomenal character like ours). And it seems that our convictions on this score are grounded in our sensitivity to the natures of phenomenal-squares and phenomenal-straight-lines.

⁴² There is a recent tradition of using what Siegel has called 'phenomenal contrasts' in order to argue for the existence of certain types of phenomenal contents (Siegel: 2010, Bayne: 2009, Green: 2017). Claim 2 explains why I was precluded from using such a method in arguing for the existence of phenomenal-composition.

I now introduce a third point in favour of claim 2:

- c) If we assume not only that claim 2 is true, but also that we are introspectively aware of the truth of claim 2 (as a) above claims), then this helps us to understand how and why we think of spatial composition as a metaphysically necessary relation.

In chapter 6 I will consider the various possible sources of our knowledge of the *metaphysical necessity* of spatial composition relations (chief among them, that it is an analytic truth that squares are composed of straight lines). There I will argue that the most plausible explanation of this knowledge does involve the postulation of perceptual experience of the metaphysical necessity of spatial composition. My concern in the current chapter is with phenomenology, not epistemology. In this section my primary concern is to argue that phenomenal-composition is a metaphysically necessary relation, so I will wait until chapter 6 to offer a complete argument for the claim that, in virtue of the necessity of phenomenal composition, we perceptually experience the metaphysical necessity of spatial composition. For this reason, I will here limit my discussion of claim c) to just those points that serve to *elucidate* that claim, and which serve to thereby highlight the relevant features of our visual *phenomenology*, which are mentioned in claim a).

So, our concern in this section really is with the phenomenal character, but if the discussion of the transparency of experience has revealed any insights at all then it is that we often find that in order to attend to features of phenomenal character we must attend to what we seem to perceptually experience (this seems to be an implication of Nida-Rumelin's (2007) Transparency Claim 6 – see the Introduction). In the present context we must attend to the way in which, as I shall contend, we seem to perceptually experience the metaphysical necessity of spatial composition. First I shall elucidate claim c).

Clearly we do think that spatial composition relations are necessary relations. But could we explain this way of thinking about spatial composition relations in the same manner that we explain our way of thinking about other necessary relations? For

example, could we explain our way of thinking about spatial composition relations in the same manner that we explain our thinking about natural laws? Though some people do believe that we perceptually experience causal connections, we can also explain why we think of causal connections as necessary just by appealing to our capacity to perform inductive reasoning.

If we were to explain our thinking about spatial composition relations just in terms of our capacity to perform inductive reasoning then we would have to say something along the following lines:

We believe that squares are composed of straight lines because every time anyone has ever seen a square, they have seen it to be composed of 4 straight lines. It is because no one *ever* sees a square not composed of 4 straight lines that we believe that squares are necessarily composed of straight lines.

Now, if our observations regarding the phenomenal-unity of phenomenal-squares and phenomenal-straight-lines are correct then there is reason to believe that inductive reasoning of the kind described above may not be necessary in order for us to recognise the necessity of spatial composition relations. In section 2.1 I said:

‘Visual experience is not only sensitive to the existence of the composition relation, but it is also sensitive to the profound intimacy of that relation. But in order that we might be perceptually sensitive to the intimacy of that relation it is crucial that the phenomenal relation between phenomenal squares and phenomenal-straight-lines be correspondingly intimate – this phenomenal unity/intimacy just *is* the perceptual manifestation of the corresponding mind-independent unity/intimacy’.

If that is correct then we can come to recognise the necessity of spatial composition relations simply by taking our perceptual experience at face value, and this requires only a single such experience. But I now wish to suggest that inductive reasoning not only is not necessary for such an understanding, but it also is not sufficient. This is

because we do not think of the necessity of the composition relation as being of the same kind as other, empirically established, generalisations.

Most people assume that the causal laws governing causal relations are merely empirically necessary, so there are possible worlds in which the causal laws fail and the same types of events can be otherwise causally related. But in the case of spatial composition we think that the necessity of the relation is *metaphysical*, so there is no possible world in which the same spatial properties can be otherwise compositionally related to one another (or not compositionally related to one another at all). If someone could not grasp the necessity of a composition relation on the basis of looking at one example of it, it is dubious that more examples would help. But what is to explain this difference in how we conceive of the necessity of causal relations and composition relations?

I propose that the reason we think of spatial composition as a *metaphysical* necessity is because phenomenal-squares arise from the intrinsic nature of phenomenal-straight-lines, and (as a) suggests) we are sensitive to this fact. Through phenomenal-composition we are perceptually aware not only that squares are *somehow* related to straight lines, we are aware of the particular *way* in which squares are related to straight lines; that is, we are aware that the square arises from the intrinsic nature of the 4 located straight lines. When this claim is stated in this positive way it can seem baroque. But the positive claim is really just the corollary of the following very plausible negative point: phenomenal-composition does not seem to be anything extra that is added to our experience of the square and the located straight lines (and which might be imaginatively subtracted from those experiences, while leaving the experiences of the composed and composing shapes in tact). It is through this feature of our perceptual experience that we are perceptually aware of a metaphysically necessary relation *as* a metaphysically necessary relation.

The implication of Claim 2 is that there could not be a subject who experienced squares and straight lines in exactly the *same way* that we do without the special kind of unity among square experiences and straight line experiences that I have described. But to say this, in and of itself, does not preclude the possibility of someone

with an alien kind of experience, who experiences squares and straight lines, but who does not even seem to experience any composition relation between them. Or someone who might even have an experience as of a square, but not an experience as of any straight lines. However, the experience of such a subject would not involve the phenomenal properties that we have designated with the terms 'phenomenal-square' and 'phenomenal-straight-lines'. For these phenomenal properties do necessitate one another. So the phenomenal character of their experience would have to be entirely alien⁴³.

Reflection on our own phenomenology simply suggests that *our* experience is not of a kind such that we could have experience as of a square and not at the same times as of any straight lines. Likewise, and for the very same reason, reflection on our own phenomenology suggests that our experience is not of a kind such that we could have an experience as of a co-located square and set of straight lines, but not as of any composition relation between them (though we certainly might fail to attend to the composition relation).

And now we are in a position to see why phenomenal-composition cannot be the product of cognitive penetration. If phenomenal-composition were the product of cognitive penetration then it would be possible to perceptually experience shape properties without perceptually experiencing the composition relations that obtain between them (that is, if we did not have those background beliefs that ostensibly penetrate the perceptual experience). But the observation that we've made is that this does not seem to be possible. Phenomenal-composition arises from the intrinsic nature of phenomenal-squares and phenomenal-straight-lines. It's not something extra that might be missing if we just so happened to lack certain background beliefs.

⁴³ Whether or not such alien perceptual experience is actually possible is something that I will remain neutral on. If, for example, we are successful in establishing that *our* perceptual experience is as the Naïve Realist describes it, then there remains a further question as to whether or not this is a metaphysical necessity; that is, there remains a further question as to whether or not *all* perceptual experience must be like *that* (that is, Naïve Realist). The example of alien perceptual experience described above is at least an epistemic possibility.

Let's turn to claim 4. If the necessity of phenomenal-composition were the product of cognitive penetration then presumably it would be the result of the belief that the composed shape – the square – was *necessarily* composed of the composing shapes – the straight lines. But if the necessity of phenomenal-composition were the product of cognitive penetration then it would be possible for phenomenal-composition to *fail* to be a necessary relation, should one not possess the relevant belief. For example, if one did not possess the background belief that the composed shape – the square – was *necessarily* composed of the composing shapes – the straight lines – then, by hypothesis, the corresponding instance of phenomenal-composition would fail to be a necessary relation.

Another way of putting the above is that, if the necessity of phenomenal-composition were the product of cognitive penetration, then the relevant belief would be the *cause* of the necessity of the instance of phenomenal-composition. But something cannot be caused to be metaphysically necessary. At least, something cannot be caused to be metaphysically necessary by any *contingent* spatio-temporal event or entity⁴⁴. To say that a relation can be *caused* to be metaphysically necessary is somewhat like saying that something non-living is non-living precisely because it was born. A birth, by its very nature, produces a life. Metaphysical necessity, by its very nature, is uncaused.

With several arguments in support of Claim 2 in place, with the entailment of claim 3 by claim 2 explained, and with the motivation for claim 4 explicated, I now turn, in the final section of chapter 2, to a potential counter-example to Claim 2.


4.4. Kanizsa Shapes: A Counter-Example to Claim Two?

In this final section of the chapter I consider an objection to Claim 2.

It may be objected that the Kanizsa Square Illusion is a counter-example to my claim that phenomenal-squares necessitate phenomenal-straight-lines (and phenomenal-

⁴⁴ I will prescind from issues around the causal efficacy of God in relation to the metaphysically necessary.

composition relations)⁴⁵. This would be a counter-example if one found it compelling that there really is a phenomenal-square (that is, a hallucination of a square), but that there are no phenomenal-straight-lines (that is, no hallucination of straight lines).



7.

Kanizsa Square.

I'm unsure how widespread this intuition is. The intuition certainly is not shared by Moore, Yantis and Vaughan (1998). They say the following concerning Kanizsa shapes:

'This is called *modal* completion because the perception of the completed surface is phenomenally experienced; that is, the edges between the inducing regions—edges that physically do not exist—produce an explicit sensory experience'. (1998: 105).

Given that the 'edges' in question are straight (or would be, if they existed), Moore et al.'s claim that these edges 'produce an explicit sensory experience', and that they are 'phenomenally experienced', implies that the experience does involve phenomenal-straight-lines (where there are no real straight lines). On the same page they also say that in this case perceptual completion produces 'illusory contours' (which illusory counters would again be straight lines). Similarly, Nanay says of perception of

⁴⁵ Thank you, Tom Stoneham, for suggesting this objection.

Kanizsa shapes that ‘we experience contours that are not present in the figure we are looking at’ (2009, 245)⁴⁶.

If Moore et al. were correct then the Kanizsa square would not be a counter-example to claim two, for the experience would not be an experience involving a phenomenal-square without any phenomenal-straight-lines. I agree with Moore et al. that the Kanizsa square involves no such experience, but I demur on the reason why. I think that there is an important difference in the way in which we perceive the non-illusory contours – those that define the genuinely present and visible corners of the Kanizsa shape - and the way in which we perceive the illusory contours (those that, in some sense, seem to connect the four corners). And we need to do justice to this difference in the way in which we perceive the illusory and non-illusory contours. If we say, as Moore et al. do, that the edges ‘produce an explicit sensory experience’ then we will find ourselves unable to do justice to this distinction. Consequently, we would be unable to explain why a perceptual experience of a Kanizsa square is any different to a perceptual experience of an ordinary square.

Moore et al. describe Kanizsa shapes as cases that resemble those cases in which ‘the lighting within a scene... render(s) real object boundaries invisible’. This could be because spotlights are being shone on the four corners of the square, but the rest of the square is unilluminated, or (as appears to be the case in the picture above) because the central portions of the square are against a background the same colour, but the corners are not (so the light reflected from the central portions of the square is the same as the light reflected by the background). But when it comes to Kanizsa shapes there are no ‘real object boundaries’ and lighting doesn’t really have any role to play at all, it just, in some sense, *seems* as though there are real boundaries, and as

⁴⁶ Nanay’s (2009) statement does not imply that the occluded straight lines feature in phenomenal content of *perception*, because he believes that the experience is one of mental imagery, not perception. But even on the assumption that the experience is perceptual, it may not quite imply that they do feature in the *phenomenal* content of perception. This depends upon whether or not it is possible to in any way perceptually *experience* a property without that property featuring in the phenomenal content of perception. This is a question to which I shall return later in this section. So far I have argued (in chapter 1) that this is not generally so, but as we shall see, there are particular (though certainly not infrequent) circumstances in which it may be so.

though light has rendered them invisible. We must account for the sense in which the boundaries seem to exist, but, crucially, we must also account for the sense in which the illusory boundaries are, at the same time, manifestly 'invisible'. If we say that all of the contours – both illusory and non-illusory - are 'phenomenally experienced' (as Moore et al. do), then it becomes impossible to account for the sense in which the illusory boundaries positively appear to be 'invisible'.

This is a perplexing problem, but I will now venture a tentative proposal. I am inclined to think that there is neither a phenomenal-square nor (the relevant) phenomenal-straight-lines. But if we say this then we need to give an account of the *sense* in which we do seem to perceptually experience a square⁴⁷. The explanation that I will offer draws on Huang and Pashler's (2007), and Huang's (2010), model of visual attention, which I made use of in chapter 2. I will utilize Campbell's (2014) interpretation Huang and Pashler's distinction, between selection and access, to explain exactly how the property type (square, triangle, or whatever) might feature in perceptual experience, without implicating a corresponding phenomenal type⁴⁸.

I wish to suggest that in the case of the Kanizsa Square we are not phenomenally conscious of any square, but we do have visual access to a (in this case, merely depicted) square. Recall, Campbell describes the distinction between selection and access thus:

'In visually attending to a scene, one dimension of your experience has to do with the characteristics of the objects that you would report them to have, act with respect to, or report yourself as experiencing. But another, more fundamental dimension of visual experience has to do with how you grab the

⁴⁷ Louise Richardson has suggested to me a 'doxastic/cognitive account of the illusion'. The problem with this suggestion is that it does not seem to locate the illusion in the right place - perception. In fact, it renders the phenomenon a kind of delusion, rather than an illusion. I want to offer an account of the illusion on which it counts as a genuinely *perceptual* phenomenon.

⁴⁸ Given that I believe that properties may, in this way, come into perceptual experience without implicating any corresponding phenomenal types, I remain open to the idea that there is a kind of content, attributable to perceptual *experience*, other than phenomenal content.

object in the first place; how, in vision, you snatch it out from the rest of the visual array as something on which you are going to focus' (Campbell, 2014: 51).

Campbell is here defending the common sense view of visual experience, on which we are simultaneously presented with a rich manifold, or array of entities, in our immediate environment. This manifold, or array, which we standardly call the 'visual field', is defined by a spatial region, relative to the subject's eyes, the contents of which are *simultaneously* visible to the subject. On this view, the visual field presents us with a range of entities that we may then go on to single out for closer inspection. But since the content of the visual field is independent of, and determined prior to, what the subject actually goes on to inspect, it is a mistake to assume that only what is reportable is conscious (i.e. part of the subject's visual experience).

As we saw, this common sense view can be given a very precise formulation in the terms of the selection-access distinction. In these terms, a property enters into the *phenomenal content* of perceptual experience when that property is available (because it falls within the visual field) for us to use to *select* an object as figure from ground. When we thus *select* an object as figure from ground, we are then able to inspect, or *access*, that object's properties.

What is pertinent to the Kanizsa case is that it seems that even when parts of the facing surface of an object are not visible (because of lighting or occlusion), still the visual system is able to select the object. My proposal is that in such cases it is only the properties of those visible (appropriately lit, and unoccluded) portions of the object's facing surface that are available for use in selecting the object as figure from ground, but that this often suffices for one to successfully select the whole object. Since only the properties of the visible (appropriately lit, and unoccluded) portions of the object's facing surface are available for use in selecting the object as figure from ground, only the properties of those portions of the facing surface feature in the *phenomenal content* of the experience. Nonetheless, once the object has been selected, we may then have visual access to properties of the whole facing surface (and even to properties of the whole, 3-dimensional, object).

Moore, Yantis, & Vaughan distinguish between two relevant types of case:

‘For example, an object can be partly occluded by other objects or by parts of itself... this is called *amodal* completion because the completed surface does not produce an explicit sensory experience. In addition, the lighting within a scene can create edges that correspond to no object boundary (e.g., shadows), or can accidentally render real object boundaries invisible... This is called *modal* completion because the perception of the completed surface is phenomenally experienced; that is, the edges between the inducing regions—edges that physically do not exist—produce an explicit sensory experience’ (1998: 105).

As I have said, it is possible to deny that the latter example involves ‘phenomenal experience’ of the illusory contours, and Moore, Yantis, & Vaughan’s reference to ‘invisible boundaries’ highlights why we might want to deny this. If we ‘phenomenally experience’ the invisible boundaries, in what sense do they present themselves as ‘invisible’? What is important is that Moore, Yantis, & Vaughan have established that we are able to use the properties of the facing surfaces of an object to *select* the whole object as figure from ground, even when part of the facing surface is occluded (this is the occluded object case), or when ‘the lighting within a scene... render(s) real object boundaries invisible’ (this is the Kanizsa shape case). And not only have they demonstrated that we are able to do this, but they have demonstrated that we *automatically* do so, even when this interferes with the performance of a task (Moore, Yantis, & Vaughan: 1998).

The mechanisms involved in our perceptual experience of Kanizsa shapes are precisely the same as those involved in our perceptual experience of occluded surfaces. The only difference is that in the former case one does not perceptually experience the whole facing surface because another opaque surface blocks one’s view, but in the latter case the visual system is working on the assumption that one does not perceptually experience the whole facing surface because of ‘the lighting within a scene’. Of course, in the case of Kanizsa Shapes this assumption is false,

because Kanizsa shapes are illusions, but this is an irrelevant feature of the situation. Below is an image of the type of display that Moore, Yantis, & Vaughan (1998) used for their occluded object condition. This display helps to highlight just how similar the two cases are, for in this display, no less than in the Kanizsa display above, there really are no contours joining what seem to be different parts of a single occluded object.




Image used by Moore, Yantis, & Vaughan (1998).

It's difficult to see why we should count Kanizsa shapes as illusions, but not also count displays like b, above, as illusions (or 2-dimensional pictures, of 3-dimensional objects, in general, for that matter). But the key point is that, for present purposes, the putative illusory dimension of the situation is irrelevant.

Once we see that the mechanisms at work are the same as those that are at work when there is an object with a square facing surface, but this surface is either partially occluded, or unevenly illuminated, we can see that the Kanizsa 'illusion' is really just the result of the fact that our visual access to an object's properties goes beyond what is phenomenally conscious⁴⁹. The contours (or edges), real or illusory, seem invisible because they are not *phenomenally* conscious (that is, there are no elements of phenomenal character corresponding to them). The contours are merely visually accessible.

⁴⁹ So, the question (raised in footnote 48) of whether or not perceptual *experience* can have content other than its phenomenal content turns (at least in part) on whether or not we wish to say that perceptual experience affords the subject with non-inferential warrant for judgments concerning those properties which are *merely* visually accessible. And it does at least seem *prima facie* plausible that perceptual experience affords such warrant.

So, we can account for the *sense* in which both the square and the straight lines enter into the perceptual experience as follows: we use the properties of the parts of the object that we can see (the appropriately lit, and unoccluded portion of the facing surface of the object), which are phenomenally conscious, to *select* the object as figure from ground. We may then access the properties of this object. The Kanizsa shape enters into perceptual experience at the level of access, and not at the level of selection. The invisible contours also enter into perceptual experience at the level of access, because they enter into perceptual experience as the contours of a visually accessible shape. In this way we can make sense of the respect in which the square is present in perceptual experience while respecting the sense in which the illusory contours also perceptually appear to be invisible⁵⁰.

It ought to be possible, in principle, to put one of the claims involved in this explanation to empirical test. In chapter 2 I said that the subject could use a feature value – say, yellow, or square – to select the regions of the visual field occupied by that type of property, and features of other kinds (from other property dimensions) will be automatically added to the Boolean map. So, in principle, it ought to be possible to test whether or not subjects can use partially occluded features to generate Boolean maps. If they can then the explanation that I have offered must be false. But even if this should prove to be so, it is far from clear that Kanizsa shapes, and modal/amodal perception, count against the theses advanced in this chapter.

Before wrapping up this final section of chapter 2 I should address one last objection to my account of the Kanizsa Square. The objection is that if perceptual access, in my

⁵⁰ Campbell (2014) says: ‘grasp of colour and shape concepts depends on the ability to access colour and shape in perception; but it seems to matter, in these cases, that we are accessing colours and shapes that figure in experience’ (66). I am suggesting that there might be visual access to a property that is not phenomenally conscious. Does this conflict with Campbell’s contention? No, it does not. For Campbell does not deny that access is ever possible without phenomenal consciousness of the relevant property, and blindsight may be a case in point. Campbell is talking about the importance of phenomenal consciousness for explaining our grasp of concepts, and it is plausible that someone who was only ever presented with *Kanizsa* ‘squares’ would not be capable of grasping the concept of a square. However, see chapter 5 for some reasons for doubting Campbell’s criterion for concept possession.

view, suffices to get properties into a perceptual experience, why shouldn't we offer the same account of our perceptual experience of composition relations? In other words, why should we think that spatial composition relations feature in the *phenomenal* content of perceptual experience?

There are two ways this possibility might be worked out. This might be because there are both phenomenal-wholes and phenomenal-parts, but the composition relation between them merely enters into perceptual experience as an accessible property. Alternatively it might be that there are only phenomenal-parts, and both the spatial wholes and the composition relations merely enter into perceptual experience as visually accessible properties⁵¹. Either way, the implication would be that it is possible to accept that we perceptually experience composition relations without accepting that there is phenomenal composition (That is, to deny the claim that in section 2 I called 'P').

I'll first address the former version of this view, on which there are phenomenal-wholes, but no phenomenal-composition. Like the advocate of the seeing-that only view, the proponent of this objection would also be forced to maintain either that there are no similarity relations among phenomenal types that mirror the similarity relations among corresponding spatial properties, or that any such similarity relations are epistemically redundant. These points proved decisive against the seeing-that only view, and there's no reason to think they prove anything but decisive in the present case.

The above response has no traction against the view on which *both* spatial wholes and composition relations are merely visually accessible. For if spatial wholes are merely visually accessible then there are no phenomenal-wholes, and the question whether phenomenal wholes bear similarity relations to one another never arises. But we have some further resources for refuting views of this kind. This second response, which I offer below, applies to both versions of the view under consideration.

⁵¹ In order that there be no phenomenal-wholes at all, the phenomenal-parts would have to be atomic. I won't pause to question the cogency of this idea.

While perceptually experiencing the Kanizsa Square (in whatever sense we do) does seem somehow *irresistible*, we are not tempted think that such an experience is *metaphysically* necessitated by the experience of the four ‘corners’. It seems perfectly possible that there might be a creature, with phenomenal character just like our own, which was incapable of either modal or amodal completion. Such a creature would see no Kanizsa Square, though the phenomenal character of their experience of the four ‘corners’ would be just like our own.

As I argued in sections two and three, matters are different when it comes to experiences of ordinary, unoccluded, spatial wholes. In such cases, assuming that the phenomenal character is like our own, it does seem compelling that experience of the parts metaphysically necessitates experience of the whole, and *visa versa* (though one may, of course, fail to attend to either of these features of one’s experience). The explanation that I offered was that phenomenal-composition has its source in the intrinsic natures of the phenomenal-whole and the phenomenal-parts. But if there were no phenomenal-composition, and if composition relations were merely a visually accessible feature of the scene, then we would be robbed of the explanation that we offered of this conviction.

4.5. Summary

In this chapter we established the following 4 claims:

- 1) Spatial composition features in phenomenal content and (consistent with Thesis 3), there is a corresponding element of phenomenal character, which I shall call ‘phenomenal-composition’.
- 2) Phenomenal-composition is a metaphysically necessary relation. It arises from the intrinsic nature of the composed and the composing phenomenal-shapes. The fact that phenomenal-composition arises from the intrinsic nature of the composed and the composing phenomenal-shapes explains why it is so easy to miss (and why it so far *has* been missed). Though it is a distinct element of

phenomenal character, it is not something *extra* in the sense that it could be imaginatively subtracted from the composed and the composing phenomenal-shapes.

- 3) Phenomenal-composition is not the product of cognitive penetration.
- 4) The *necessity* of phenomenal-composition is not the product of cognitive penetration.

Chapter 5 – Knowledge of Spatial Composition.

In chapter 4 we considered an alternative to perceptual experience of spatial composition – the seeing-*that* only view. This alternative requires that we have general beliefs about the relationship between composed and composing shape properties, which we did not acquire through perceptual experience of those relationships. For example, the view requires that we believe that whenever a square is co-located with 4 straight lines then the square is composed of those straight lines, and it requires that we did not acquire this belief through perceptual experience. In this chapter I consider where such knowledge might come from.

One obvious account of our knowledge of such general truths regarding composition relations is that they are analytic truths. Analytic truths are statements that are true in virtue of their meaning. The implication is that one can recognise such statements as true provided only that one understands them. In section 1 of this chapter I consider how the proponent of the seeing-*that* only view might appeal to analyticity, and I will consider some very general problems with the idea of analyticity.

In section 2 we shall consider a model of concepts – the classical model - that would explain how certain truths might be analytic, and I will argue that even if this model were correct, still those truths pertaining to spatial composition relations would not be analytic. Then, in section 3, we shall consider what appears to be the only remaining option available to the proponent of seeing-*that* only view: that we have innate beliefs concerning spatial composition relations. This response does provide the proponent of the seeing-*that* only view with a cogent answer, but it forces them into a fresh confrontation with the phenomenological observations made in the previous chapter.

Nothing that I say here will amount to a knock down argument against the seeing-*that* only view, but, I believe that the phenomenological case that I made against this view in the previous chapter is already strong, and an examination, in this chapter, of

some of the implications of the view will serve to further strengthen the argument against the view.

5.1 Analytic Truths

5.1.1 The Relevance of Analyticity.

One suggestion that can be immediately ruled out is that the source of the knowledge is testimony. This answer might be true on the individual level, but it would leave it mysterious where the knowledge ultimately originates from. In this regard, to say 'testimony' is simply to defer the question. It leaves us none the wiser as to how society at large acquired knowledge of spatial composition relations. Arguing in this way means that there is leeway for an opponent to concede that my argument establishes that *someone* must have once perceived spatial composition relations, but to maintain - in the face of the phenomenological argument that I propounded in chapter 4 - that *our* knowledge of spatial composition relations is testimonial. But I think that this position would decrease the plausibility of the seeing-that view to such an extent that I am unconcerned by the logical coherence of this line of argument. If we have reason to believe that one of us has perceptually experienced spatial composition relations then we have good reason to believe that we all do (especially given the phenomenological argument advanced in chapter 4)⁵².

For those opposed to the idea that we perceive spatial composition relations, an obvious explanation of the origin of our beliefs pertaining to such relations is that these beliefs concern analytic truths. That is, that the truths are somehow 'inherent' in the meanings of the concepts involved. The reason that it might help the advocate of the seeing-*that* only view if the relevant truths are analytic is because it is usually assumed that analytic truths are *a priori* truths. If it were an *a priori* truth that

⁵² Similarly, my argument leaves untouched the logical coherence of the view that someone learnt of spatial composition relations through divine revelation, and the rest of us know of them only through testimony (or even that all of us know of them through divine revelation, but that most of us are somehow ignorant of this fact).

squares are composed of straight lines then it would be possible to understand how we have knowledge of such truths without appealing to perceptual experience.

Now, there are legitimate reasons for which one might propose an '*a posteriori* analytic' category of truths. If one assumes that 'meanings aint in the head', there is no reason to think that some truth that is somehow inherent in the meanings of the concepts involved should also be an *a priori* truth. For example, if one holds that names are rigid designators – i.e. their meanings are exhausted by their referents – then one might argue that it *is* inherent in the meaning of 'the morning star' that it is identical with the evening star. After all, the meaning of 'the morning star' just is the object that we now call 'Venus', and it is necessary that Venus is identical with itself, no matter what name we use to identify it. So, it might reasonably be supposed that it is inherent in the meaning of 'the morning star' – that is, the object, Venus – that it is identical with The Evening Star – that is, itself.

The possibility of *a posteriori* analytic truths will not help the proponent of the seeing-that view. For if *these* analytic truths are *a posteriori* then there remains a role for perceptual experience in the acquisition and justification of such knowledge. For example, observations of certain kinds were crucial in the discovery that the morning star is identical with the evening star (and in the justification of this claim). The trouble is that meaning is not always entirely transparent to us, even where we do grasp the relevant thought, or understand the relevant sentence. The proponent of the seeing-that only view needs it to be the case that:

- 1) There is a class of truths that one may recognise as true simply in virtue of understanding them (i.e. a class of statements that are justified by their meaning, where their meaning is transparent).
- 2) Truths pertaining to spatial composition relations belong to that class.

The proponent of the seeing-that only view wishes to dispense with any role for perceptual experience in our knowledge of spatial composition, so they require what Williamson calls an *epistemological* conception of analyticity. A conception of analyticity that would serve their needs would be one on which 'linguistic or

conceptual competence constrains one's attitudes to analytic sentences or thoughts' (Williamson, 2007: 73). Clearly, mere linguistic or conceptual competence will not suffice to constrain one's attitude towards the sentence 'the morning star is identical with the evening star'. The meaning of that sentence is not entirely transparent, even for one who does understand the sentence. One can understand the names, and understand the sentence that they appear in, and yet, in so doing, one may have no inclination or reason to assent to the sentence. But might mere linguistic or conceptual competence somehow suffice to constrain one's attitude towards the sentence 'squares are composed of straight lines'?

In section 1.2 I will appeal to Williamson's assault on 'epistemological conceptions' of analyticity, in order to cast some doubt over 1) above. Then in section 2 I argue that, whether or not 1) is true, 2) is false.

5.1.2. Williamson, and General Problems with the Epistemological Notion of 'Analytic Truths'.

Williamson makes the following, *prima facie* plausible, observation:

'If someone is unwilling to assent to the sentence "every vixen is a female fox," the obvious hypothesis is that they do not understand it, perhaps because they do not understand the word "vixen". The central idea behind epistemological conceptions of analyticity is that, in such cases, failure to assent is not merely *good evidence* of failure to understand; it is *constitutive* of such failure (2007: 73).

Yet, early in his chapter concerning epistemological conceptions of analyticity, Williamson says 'in what follows, we will consider more rigorously what is epistemically available simply on the basis of linguistic and conceptual competence. To a first approximation, the answer is: nothing' (2007: 77).

He suggests that what epistemological conceptions of analyticity require is what he calls “understanding-assent links”. He offers the following examples of prima facie plausible candidates:

‘(UAI) Necessarily, whoever understands the sentence “every vixen is a female fox” assents to it’ (2007: 73).

‘(UAt) Necessarily, whoever grasps the thought *every vixen is a female fox* assents to it’ (2007: 74).

An initial reservation that one might feel is that it seems that one can understand the above sentences without explicitly assenting to them - either overtly or covertly. But Williamson notes that assent is dispositional. One need not have actually assented to the relevant sentence in order to understand it. Nonetheless, one may still feel some resistance to the idea that understanding-assent links are involved wherever there is an analytic truth. Where analytic truths are very complicated, arguably, one may understand them without even being disposed to assent to them (complicated mathematical truths provide a good example). However, if one is to hold onto the sense in which the meaning of a sentence might actually serve to constrain one’s attitude towards the sentence, then rather than entirely jettisoning the understanding-assent links, one ought to instead suitably weaken them. Williamson proposes to start by examining the unqualified understanding-assent links, and then to later consider how they might be loosened.

Of course, understanding-assent links alone will not suffice for an epistemological conception of analytic truths, for we require *knowledge* of analytic truths. To be more precise, we require knowledge of analytic truths, where the relevant beliefs are *justified* by the meaning of the relevant propositions – not, for example, by testimony. I will call such knowledge ‘analytic knowledge’. In this regard, Williamson says:

‘On the most optimistic view, understanding-assent links generate understanding-knowledge links like these:

(UKl) Necessarily, whoever understands the sentence “every vixen is a female fox” knows “every vixen is a female fox”.

(UKt) Necessarily, whoever grasps the thought *every vixen is a female fox* knows *every vixen is a female fox*. (2007: 77)

The idea of an analytic truth is that one can know the truth merely in virtue of understanding the words (by grasping the corresponding concepts) involved. For example, if it is a condition on understanding the word ‘square’ that one know the following definition: ‘shape composed of 4 straight lines of equal length’, then one can know the truth of the statement that ‘every square is composed of straight lines’ simply in virtue of understanding the words, and the modes of combination, involved in the statement.

Generalising, the proposal is that for each word there are a number of truths that one can know simply in virtue of understanding that word. This entails that for each word there are a number of truths that one must assent to (or be disposed to assent to) if one is to count as understanding that word. And since knowledge requires assent, it also requires that when one asserts or judges these truths on the basis of one’s understanding of the terms involved, one thereby knows (i.e. minimally, one is justified in believing) their truth. Williamson notes that there are complications in getting from understanding-assent links to understanding-knowledge links. But his main focus is on the question of whether there even are any understanding-assent links. For if there are no understanding-assent links then the project will be forestalled before the question of understanding-knowledge links even arises.

Quine (1951) famously challenged the analytic/synthetic distinction by means of attacking the notion of ‘synonymy’. Unlike Quine, Williamson accepts that ‘synonymy’ is a perfectly legitimate notion: ‘By ordinary working standards, the word “synonymous” is quite clear enough to be useful’ (2007: 50). He holds that two words are synonymous if and only if they have the same *intension* (i.e. function from context to extension). So, with this notion in hand, the advocate of analytic truths may maintain that ‘square’ is synonymous with ‘shape composed of 4 straight lines of

equal length'. For Williamson the issue, rather, is that synonymy is not entirely transparent to just anyone who understands the relevant pair of terms. This is because the intension of a term is not entirely transparent to just anyone who understands that term. It is possible to understand two phrases, which are synonyms, – 'vixen' and 'female fox', say - without knowing that those terms are synonyms. And someone who does understand two terms, without knowing that they are synonyms, may not even be disposed to assent to the corresponding putative analytic truth – i.e. the one that trades on their synonymy; in this case 'every vixen is a female fox'.

So we have the following putative analytic truth:

1. Every vixen is a female fox.

But Williamson observes that 'someone under the misapprehension that the term "vixen" also applies to immature male foxes may believe that every vixen is a vixen without believing that every vixen is a female fox' (2007: 117).

Alternatively, Williamson suggests:

'We can imagine that our speaker is quite familiar with the dictionary definition of "vixen" as "female fox." He also knows that dictionaries give a second definition of "vixen" as "quarrelsome woman." However, unlike most of us, he does not believe that these are two senses of "vixen". Rather, he thinks that "vixen" in its primary sense applies to both female foxes and quarrelsome women. He may defend his view with sophisticated arguments from the philosophy of language, although this is not essential' (2007: 118).

Williamson goes on to say:

'Our imaginary speaker is not so different from actual native speakers of English who deny that a man who has lived with a partner for several years without getting married is a bachelor, or assert that someone who underwent

a sex-change operation after giving birth is a mother without being a female parent'. (2007: 118)

Unless one denies that the envisaged speakers understand the words 'vixen', 'bachelor,' and 'mother', these examples demonstrate that synonymy is not transparent; it is possible to understand two synonyms but to not know that they are synonyms. This would mean that we have a counter-example to an understanding-assent link for 1 above. But is it true that such speakers do understand the words 'vixen', 'bachelor,' and 'mother'?

If the advocate of an epistemological conception of analyticity denies that such speakers understand the words 'vixen', 'bachelor,' and 'mother' then they can also deny that the relevant speakers understand 1 above. In which case, these speakers do not count as counterexamples to the understanding-assent links (UKl) and (UKt). On this picture, presumably these speakers use the relevant words in an idiosyncratic way. The thought that they refuse to assent to in refusing to assent to 1 will then be a different thought to that which we associate with 1. And presumably the thought that they refuse to assent to, unlike that which we associate with 1, is not an analytic truth.

The assumption behind the above picture seems to be that one must have perfect knowledge of the meaning of some word in order to count as understanding that word. But Williamson finds this implausible. He remarks: 'One can know that "red" means *red* without being infallible as to exactly which shades count as shades of red' (2007: 124). If someone were to deny that some peripheral shade of red were a shade of red at all - perhaps they claim that it is instead a peripheral shade of orange - then we would attribute to this person the belief that this shade is not a shade of red. But this is a belief that only someone who possesses the concept RED (and who is thereby capable of grasping thoughts about the colour red) can hold, for it is a belief about which shades are shades of the colour *red*.

Likewise, talking of those speakers who denied that a man that has lived with a partner for several years without getting married is a bachelor, Williamson says:

“They are mistaken about the meaning of the English words “bachelor” and “unmarried.” Nevertheless, they fall well within the range of permissible variation for linguistically competent speakers... We classify them as believing that some unmarried men are not bachelors and that some mothers are not female parents because we interpret them as having used the words with their normal English meanings, despite their errors. That is how they intend to be interpreted, not as using the words with idiosyncratic senses (2007: 118).

One might wonder at this point whether these hypothetical subjects really do have (false) beliefs about bachelors or about the colour red. Perhaps, one might speculate, these subjects really just have beliefs about the word ‘red’ and the word ‘bachelor’. More specifically, perhaps these subjects have false beliefs about the extensions of these words – i.e. the belief that the word ‘red’ does not apply to *this* shade, and the belief that the word ‘bachelor’ does apply to some unmarried men. Consistent with this, these hypothetical subjects might not possess the concept RED or the concept BACHELOR, on account of their ignorance of the precise extension of the concepts/corresponding words (though they would, of course, possess the concepts of the corresponding words ‘red’ and ‘bachelor’). One could then maintain that these subjects do not understand the sentences that they assert. Given their idiosyncratic meta-linguistic beliefs, we can understand why they would assert these sentences without supposing them to actually believe what the sentences express. In which case, since they do not understand the sentences, these subjects would not be counterexamples to the corresponding understanding-assent links.

But Williamson adds:

‘Arguably, their error is not primarily semantic: they have the semantic belief that the word “bachelor” does not apply to all unmarried men because they have the non-semantic belief that some unmarried men are not bachelors and the semantic knowledge that “bachelor” applies only to bachelors’ (2007: 119).

This passage sounds right, for we can imagine that speakers might assent to the statement “all unmarried men are bachelors” but then change their mind after they were presented with the case of a man who has lived with a partner for several years without getting married. So any deviant semantic beliefs here seem to be derived from intuitions about particular cases; about whether *this* unmarried man is a bachelor.

Grice and Strawson, responding to Quine, suggest that ‘my neighbour’s 3 year old child is an adult’ is ‘a sentence that we could not understand someone using with its ordinary literal meaning to make an assertion (1956)’. The implication of this, Williamson surmises, is that there is an understanding-assent link for the sentence “no three year old child is an adult”. But Williamson points out that:

‘Someone may believe that normal human beings attain physical and psychological maturity at the age of three, explaining away all the evidence to the contrary by *ad hoc* hypotheses or conspiracy theories (many three-year-olds pretend to be eighteen-year-olds in order to vote, the abnormally polluted local water slows development, and so on). However foolish those beliefs, they do not constitute linguistic incompetence. (2007: 85)

He goes on to say ‘friends of analyticity will reply that the example was badly chosen. It is therefore best to start with the most elementary examples possible’ (2007: 85). So Williamson proposes to ‘focus... on the simplest cases, since those are the ones for which understanding-assent links have the best chance: if they fail there, they fail everywhere’ (2007: 85). His contention is that understanding-assent pairs are most plausible when it comes to the logical constants: universal quantification, conjunction, disjunction etc. Williamson considers numerous such candidates for understanding-assent links, and he offers a plethora of counter-examples. We shall consider just two of the candidates, and the corresponding counter-examples, that Williamson offers.

The first of these two candidates is the following:

Necessarily, whoever understands the sentence “every vixen is a vixen” assents to it (or is disposed to do so).

And the corresponding analytic truth is

2. Every vixen is a vixen.

This case is simpler than 1. - ‘Every vixen is a female fox’ -, because 2., unlike 1., is a *formal* tautology. Even if one does not understand the word ‘vixen’, one may be in a position to tell that 2. is true on the basis of the *form* of the sentence. However, in order to recognise the sentence as true, even if one did not understand the word ‘vixen’, one would still need to understand the word ‘every’. So Williamson proposes to demonstrate that, if one has deviant beliefs concerning quantification, it is possible for one to understand 2. and refuse to assent to it. Indeed, in the examples that Williamson offers the subjects possess intelligent reasons for denying 2.

Peter and Stephen both refuse to assent to 2., but for different reasons. Peter:

‘Regards the truth of “There is at least one F” as a necessary condition for the truth of “Every F is a G” quite generally, and the falsity of “There is at least one F” as a sufficient condition for the falsity of “Every F is a G”; he takes universal quantification to be existentially committing’ (2007: 86)

Peter believes that 2. is false because:

‘He spends far too much time surfing the internet, and once came across a site devoted to propagating the view that there are no foxes, and therefore no vixens, and never have been: all the apparent evidence to the contrary has been planted by MI6, which even organises widespread fox-hallucinations, so that people will protest about fox-hunting rather than the war in Iraq’ (2007: 87)

Stephan has other qualms:

'What worries him is vagueness. He believes that borderline cases for vague terms constitute truth-value gaps. Like many truth-value gap theorists (such as Soames (1999)), he generalises classical two-valued semantics by treating the gap as a third value ("indefinite") and using Kleene's three-valued "strong tables" (1952: 334)... On Stephen's view, for "every F is a G" to be true is for the conditional "x is an F – x is a G" to be true for every value of the variable "x"; for "every F is a G" to be false is for "x is an F – x is a G" to be false for some value of "x". On his semantics, for a conditional sentence with "-" to be true is for either its antecedent to be false or its consequent to be true, and for it to be false is for its antecedent to be true and its consequent false. Stephen also believes that some clearly female evolutionary ancestors of foxes are borderline cases for "fox" and therefore for "vixen". Consequently, for such an animal as the value of "x", "x is a vixen" is neither true nor false, so the conditional "x is a vixen – x is a vixen" is also neither true nor false, by the strong Kleene table for '-'. (2007: 87)

Of course:

'Someone might insist that Peter and Stephen appear to be using the word "every" in its standard sense because they are really using it in senses very similar to, but not exactly the same as, the standard one' (2007: 89).

Perhaps, in their mouths, 2. expresses a different proposition to that which it expresses in our mouths. In their mouth 2 might not express an analytic truth.

However, Williamson retorts:

'Peter and Stephen are emphatic that they intend their words to be understood as words of our common language, with their standard English senses. They are not making unilateral declarations of linguistic independence. They use "every" and the other words in (1) as words of the public language. Each of them believes that his semantic theory is correct for English as spoken by others, not just himself, and if it turned out to be (heaven forbid!) incorrect

for English as spoken by others, it would equally turn out to be incorrect for English as spoken by himself. Giving an incorrect theory of the meaning of a word is not the same as using the word with an idiosyncratic sense- linguists who work on the semantics of natural languages often do the former without doing the latter. (2007: 89)

And he adds:

‘Nor are Peter and Stephen marginal cases of understanding: their linguistic competence is far more secure than that of young children or native speakers of other languages who are in the process of learning English... If some participants in a debate have an imperfect understanding of one of the key words in which it is conducted, they need to have its meaning explained to them before the debate can properly continue. But to stop our logical debate with Peter and Stephen in order to explain to them what the word “every” means would be irrelevant and gratuitously patronizing... the understanding they lack is logical, not semantic’ (2007: 91).

Williamson also considers the possibility that Peter and Stephen might express the same proposition as we do when they say 2, but that they might associate a different thought with that proposition (where thoughts are individuated at the level of intension, so different thoughts can be associated with the same proposition). In which case, the thought that *they* associate with 2 might not be an analytic truth. The view, then, is that Peter and Stephen mean the same by their words as we do – that we speak a shared, public language – but also that Peter and Stephen associate different thoughts with those words (and sentences) to those that we do. Williamson thinks that is an unstable position:

‘If Peter and Stephen associate (the sentence) with different thoughts from ours, should we not understand them better by translating their idiolects non-homophonically into ours? Presumably we should seek sentences other than (that sentence) that we associate with the very thoughts that they associate with (that sentence), or at least sentences we associate with thoughts more

similar to the thoughts they associate with (that sentence), and translate the dissent from (that sentence) in their mouths as dissent from those other sentences in our mouths' (2007: 114-115)

But, he goes on to say:

'To insist on applying such a non-homophonic translation scheme to them in the face of their protests would be to treat them less than fully seriously as human beings, like patients in need of old-fashioned psychiatric treatment, whose words are merely symptoms. The claim that Peter and Stephen associate (that sentence) with different thoughts from ours repackages our disagreement with them in a way that makes it sound less threatening than it really is.' (2007: 115)

The second candidate that Williamson considers is the word 'if'. He says 'It is often claimed that assent to arguments by modus ponens of the form "If A then B; A; Therefore B" is a precondition for understanding the word "if" (2007: 92).

But Vann McGee has offered the following counter-example to this rule:

'Opinion polls taken just before the 1980 election showed the Republican Ronald Reagan decisively ahead of the Democrat Jimmy Carter, with the other Republican in the race, John Anderson, a distant third. Those apprised of the poll results believed, with good reason:

If a Republican wins the election, then if it's not Reagan who wins it will be Anderson.

A Republican will win the race.

Yet they did not have reason to believe:

If it's not Reagan who wins it will be Anderson'. (2007: 92)

Williamson says:

‘Does McGee not understand the English word “if”? In conversation, he appears to understand it perfectly well... Moreover, his doubts derive from taking at face value a natural pattern of native speaker reactions to an ingeniously chosen case. If he counts as not understanding “if”, so do millions of other native speakers of English’ (2007: 94).

Williamson considers whether one might ‘invoke the division of linguistic labour (Putnam 1975), and say that making any given inference by modus ponens is a precondition only for *full* understanding of “if”, the kind of understanding characteristic of the expert rather than the layman? The trouble is that McGee *is* an expert on conditionals. He publishes on them in the best journals. He does not defer in his use of “if” to any higher authorities’. (2007: 94)

Williamson concludes that:

‘Speakers can compensate for their deviance on one point by their orthodoxy on others, their ability to predict the reactions of non-deviant speakers, their willingness in the long run to have their utterances evaluated by public standards. As we have seen, such compensation is often possible when the deviance results from localised interference in the normal practice of using a word by high-level theoretical concerns. Thus there is no litmus test for understanding. Whatever local test is proposed, someone could fail it and still do well enough elsewhere with the word to count as understanding it. (2007: 97)

At the beginning of the section I said that it is plausible that where analytic truths are very complicated one may understand them without even being disposed to assent to them. And I said that in order to accommodate this intuition it might be necessary to weaken the understanding-assent links. But in light of what has emerged, it seems that these links will, at the very least, need to be weakened in order to deal with counterexamples to even the very simple ‘analytic’ truths that we’ve just be considering. How might this weakening of the understanding-assent links go?

Williamson considers that the advocate of an epistemological conception of analyticity might accept all of the above but argue that 'if the deviance results only from erroneous theorizing that overlays an ordinary understanding of the terms, may not the links still hold at the underlying level?' (2007: 99). Williamson notes that one might draw an analogy here with the distinction, in linguistics/developmental psychology, between 'competence' and 'performance' (there is evidence that children's linguistic 'competence' outstrips their linguistic 'performance').

The proposal, from Williamson's opponent, is that we need only somehow weaken the understanding-assent links so as to accommodate the possibility of individuals who refuse to assent as a result of erroneous theorizing or some other interfering factor. Perhaps, then, in order to understand some analytic truth one need not actually be, all things considered, disposed to assent to it. Perhaps it suffices that one just have some sub-personal (masked/overridden) disposition to do so. If this were correct then one could then embellish the view such that these masked dispositions were masked dispositions to *know* the relevant analytic truths.

The idea under consideration is that 'the postulated dispositions are grounded in something like an unconscious reasoning module' (2007: 100). The postulated rules must be encased in some sort of psychological module 'for if they consisted only in general habits of reasoning, Peter and Stephen's earlier habits could eventually be erased by their later ones, and the disposition to assent to (1) would disappear' (2007: 102). Williamson then argues against the view that we have such a module including unconscious logic rules, though we need not concern ourselves with the details of this argument here.

Williamson's conclusion is that even for those word-truth pairs for which it is *prima facie* most plausible that understanding-assent links hold, they do not. Moreover, even if they did do so, we would still be none the wiser as to how one is *justified* in believing the truths simply in virtue of understanding them. At this point one may begin to fear that Williamson has entirely jettisoned the distinction between understanding and not understanding a word. But this is not so. The difference, for Williamson, is nuanced. Williamson believes that the reference of a word 'is fixed by

the pattern of use over the whole community' (2007: 124), and, on the difference between understanding and not understanding a word, Williamson says:

There is, of course, a distinction between understanding a word and not understanding it. One can lack understanding of a word through lack of causal interaction with the social practice of using that word, or through interaction too superficial to permit sufficiently fluent engagement with the practice. But sufficiently fluent engagement in the practice can take many forms, which have no single core of common agreement' (2007: 126).

With these general problems with the notion of 'analytic truth' in the background, I now turn my attention away from the notion of an 'analytic truth' in general and onto the specific case of truths pertaining to spatial composition relations.

5.2. Spatial Composition Relations as Analytic Truths.

5.2.1. Assent and Justification.

We return now to the specific case of truths pertaining to spatial composition relations. It would serve the purposes of the proponent of the analytic truth view if they could establish a range of understanding-assent links of the following form:

Necessarily, whoever understands the sentence "every square is composed of 4 straight lines of equal length" assents to it.

Establishing the above understanding-assent link would require them to establish that 'square' and 'shape composed of 4 straight lines of equal length' are *transparent* synonyms. In fact, if it can be established that 'square' and 'shape composed of 4 straight lines of equal length' are transparent synonyms then it won't be necessary for them to establish the above understanding-assent link.

If 'square' and 'shape composed of 4 straight lines of equal length' are transparent synonyms then the problem of justification reduces to the problem of the justification

of logical truths. And if logical truths turn out not to be analytic truths (as Williamson suggests), then the truths pertaining to spatial composition relations will not be analytic truths either, and there will be no reason to expect the above understanding-assent link to hold. It is true that if 'square' were synonymous with 'shape composed of 4 straight lines of equal length' then 'every square is a shape composed of 4 straight lines of equal length' would be synonymous with 'every square is a square'. But if Williamson is correct then neither 'every square is a square', nor 'every square is a shape composed of 4 straight lines of equal length', are analytic truths. This is for the same reasons that Williamson offered against 'every vixen is a vixen' being an analytic truth.

In the present context, the purpose of the 'analytic truth' account of our knowledge of spatial composition relations is that it affords the advocate of the seeing-that view with explanation of our knowledge of spatial composition relations, which does not require them to concede that we perceptually experience spatial composition relations. Though the above possibility violates the letter of this response to my composition argument – in that, if Williamson is correct, the truths would not be analytic -, it does not violate the spirit of that response. What matters for the proponent of the seeing-that only view is that the truths pertaining to spatial composition relations will be justified by whatever justifies these logical truths, and, presumably, the justification of these logical truths, whatever that does involve, is unlikely to involve an appeal to perceptual experience, specifically, of spatial composition relations. This means that the advocate of the seeing-that only view who wishes to pursue a response along the lines that we have been exploring need only establish that pairs of spatial phrases like 'square' and 'shape composed of 4 straight lines of equal length' are transparent synonyms. For ease of expression I will continue to assume that if pairs of spatial phrases like 'square' and 'shape composed of 4 straight lines of equal length' are transparent synonyms then the corresponding truths will be 'analytic'.

5.2.2 How Might the Understanding-Assent Links Hold? Constitutive Inter-Conceptual Relations.

It is crucial to the 'analytic truth' response that the analyticity of the truths casts some light on *how* the understanding-assent links might hold. We can see how these understanding-assent links would hold if we assume the classical view of concepts. On this view the statement that 'bachelors are unmarried men', for example, is analytic because one of the concepts involved is composed of the other concepts involved – i.e. the concept 'bachelor' is composed of the concepts 'married' and 'man'.

Applying this to the present case this means that the concept 'square' must be composed of the concepts '4', 'straight lines', and 'equal length'. If this is right then we can see how the understanding-assent links would hold (though we would still need to explain how the understanding-knowledge links hold – i.e. how are the statements justified). Another way of putting the idea is that 'interconceptual relations (are) constitutive of concept possession' (Fodor, 1998: 70)⁵³. So I will assume that it is an implication of some truth being analytic that one cannot possess one of the concepts involved - the concept of a bachelor, for example - without possessing the other concepts involved in the truth – the concepts of 'marriage' and 'men', for example.

When we consider sources of knowledge at the individual level we tend to treat the analytic truths as a common class. But when we are concerned with the ultimate origins of knowledge, to say that the knowledge is analytic is also to defer answering the question. For example, if the knowledge is implicit in our concepts, and if we acquire our concepts from our linguistic community, then this is similar to learning something through testimony. Such a view would not enable us to answer the question: how did our concepts come to accurately reflect the world? And, of course, the proponent of the seeing-that only view is precluded from appealing to perceptual experience in answer to this question.

In what follows I shall consider the view that the concepts are innate. This view has a straightforward answer to the question: how did our concepts come to accurately reflect the world? Since there is survival value in having inter-conceptual relations

⁵³ This is an idea that Fodor emphatically rejects.

that accurately reflect the world, we can appeal directly to evolution in answering this question.

Recall that Williamson considered a view on which the understanding-assent links were weakened, to permit that subjects might understand the relevant sentences without actually being disposed to assent to them. In the case of logical primitives this involved postulating a logic/reasoning module, which might ground dispositions to assent to logical truths, but which dispositions might be masked/overridden by erroneous theorising. The examples we considered were Peter and Stephen, who both understood the word 'every', but who held theories that served to mask/override their dispositions to assent to the sentence 'every vixen is a female fox'. The implication is that the relevant dispositions are, in their case, merely sub-personal.

A strength of the view that the understanding-assent links hold in virtue of inter-conceptual relations (the IC view) is that it can accommodate the idea that subjects might not actually be disposed to assent (or that they are disposed to do so, but only in a very specific situation: *after* reflection on the relevant concepts). The proponent of the seeing-that only view can say that, until we subject our concepts to some kind of process of 'conceptual analysis', the assent dispositions that are grounded in the structure of those concepts are merely sub-personal (and so, in a sense, masked/overridden). Once we subject the relevant concepts to conceptual analysis we acquire the relevant beliefs, and the dispositions become personal-level.

This view implies that the concepts of spatial composition and of various types of shape property (e.g. squares and straight lines) are innate. But we don't just have to explain our capacity to see-*that* squares are composed of straight lines, we also have to explain our capacity to see-*that* a triangle is composed out of straight lines, and to see-*that* a circle is not. Indeed, the same follows for any pair of types of shapes that we can see. The implication is that all of our shape concepts must be innate.

For the seeing-*that* only theorist who advocates innate concepts the beliefs will be a priori because the subject can acquire/justify the beliefs by 'analysing' her (innate)

concepts, and without recourse to experience. The innate concepts (IC) theorist will say that our concepts are built out of one another in a way that reflects the composition among the shapes. So, on this view, the concepts of the more composite shapes are composed out of the concepts of the component shapes.

This view faces problems in the case of complicated shapes for which we seem to have no (non-demonstrative) concepts. Even in the case of such complicated shapes we can still see-*that* such shapes are composed of others in certain ways. This is proven by our capacity to replicate such complicated shapes, if given a ball of play dough.

If the IC theorist holds a *strong* nativism (like that proposed by Fodor (1998)), on which *all* concepts are innate, they might respond that since it is possible to generate names for very complex shapes, and to learn to recognise instances of these shapes (after all, we recognise individual people by recognising the intricate shapes of their faces), then we must after all be in possession of the relevant concepts. In which case, it is open to the nativist to say that these (innate) concepts, which are not associated with any natural language words, are nonetheless activated by our perceptual experiences of complicated shapes. When the spatial concepts are activated, so too are the relevant beliefs, and these beliefs are then able to explain and justify the relevant capacities.

The problem with this response is that the subject must actually perform a conceptual analysis on the relevant concepts before she will hold the relevant beliefs, and, arguably, this would require her to already have natural language words for those concepts. Yet it seems that someone who has no words for the relevant complicated shapes can still see *that* they are composed of the shapes that they are composed of, in the ways that they are.

In the next section I will explain why the understanding-assent link does not hold for the statement 'every square is a shape composed of four straight lines of equal length'. In section 1 we considered some counterexamples, offered by Williamson, to understanding-assent links for some other plausible candidates for analytic truths.

These counterexamples involved subjects who held that two synonyms were not synonyms, and who believed that one of the terms applied in cases in which the other did not; they believed the terms to have different extensions. I will not offer a direct counterexample, of the kind considered in section 1. Instead, I will offer some examples that undermine the required application of this theory of concepts to our concepts of spatial properties and their composition. In so doing, I will undermine the theory that would explain *how* the understanding-assent link would hold in this case.

5.2.3. Undermining Inter-Conceptual Relations Among Shape Concepts.

According to the classical view of concepts, at least for many concepts, you cannot be in possession of that concept without knowing, or being able to derive, some definition for it. This structure is usually assumed to bottom out with a class of observational concepts. Since these observational concepts are not built out of more basic concepts, possession of any observational concept does not entail the possession of any further concepts. I wish to argue that if one endorses the classical view of concepts, then one ought to consider spatial concepts like 'square' as among the observational concepts.

First, I will say a little more about the observational/non-observational concept distinction. One traditional worry that has prompted the postulation of observational concepts relates to concept acquisition. If the possession of each and every concept implicates possession of further concepts then it's difficult to see how a subject might ever acquire her first concept. I think that this argument has some merit, but no nativist about concepts will be persuaded by it, so I do not intend to lean on it. For this reason, I will not attempt to make any general argument for an observational/non-observational distinction.

Instead, I will focus primarily on the example of a square, and I will argue that even if some concepts are composed of others (as the classical view of concepts suggests) then SQUARE is not. The implication is that if (contra-Williamson) the classical view of concepts is correct, SQUARE must be an observational concept. In which case, the seeing-that theorist is deprived of the explanation of the relevant,

putative, understanding-assent link. That is, there is no explanation of how, or why, *understanding* the statement 'every square is a shape composed of 4 straight lines of equal length' should presuppose any disposition to *assent* to it. So the statement cannot be analytic (or, at least, we have no explanation of how it could be). It should be clear that similar considerations would apply to other spatial concepts.

The argument, in brief, is that it is possible to possess the concept of a square without possessing the concepts of straight lines or spatial composition.

Suppose that a subject is born without certain of the relevant innate concepts.

Suppose that she has no concept of a straight line, or of spatial composition. It seems perfectly possible that such a child might have the capacity to sort objects on the basis of their shapes, such as squareness, and even have some knowledge of the dispositions that being square bestows upon an object (e.g. circular objects roll, square objects don't roll), without being in a position to define a square (even after suitable reflection on her concepts), and without even being in a position to distinguish objects whose borders do involve straight lines from those that don't. This idea shouldn't seem wild. We are all very good at recognising faces, even though we might struggle to conceptualise the face's features and the shapes of these.

Alternatively, imagine that she has the concept of a straight line, but not the concept of spatial composition. It seems that it would be possible for there to be a subject who had perceptual experience just like ours, but who, due to some kind of attentional bias, never noticed that squares were composed of straight lines. She in no way evidences possession of the concept 'spatially composes'. She might nonetheless possess the capacity to recognise and sort objects both on the basis of their shapes, and on the basis of whether or not they instantiated any straight lines. Such a subject might even have a fairly sophisticated understanding of how the shapes of objects dispose those objects to behave (e.g. round things are disposed to roll, square things are not). Given that she has never noticed that squares are composed of straight lines, she will not be able to purposefully manipulate the shapes of objects of course, nor will she have the capacity to make the kinds of similarity judgements (between *determinate properties*) that I have previously described (though she might

conceivably be able to compare such *objects*, on the basis of their shared *determinable* properties⁵⁴).

But must we deny that such a subject has the concepts of squares, circles and straight lines? It seems dogmatic to maintain that we must deny the subject these concepts, simply because she is unable to produce (even after suitable reflection on her concepts) the definition of a square which we consider canonical. The implication is that the subject can possess the concepts of squares and circles without possessing the concept of spatial composition. In which case, absent any alternative explanation, it cannot be an analytic truth that squares are *composed* of straight lines.

It also seems perfectly possible that a subject might have the capacity to recognise and sort objects on the basis of their shapes, and to make predictions about their behaviour on this basis, but *not* have the capacity to sort objects on the basis of whether or not they instantiated any straight lines. Indeed, such a subject might not in any way evidence possession of the concept 'straight line'. If this seems implausible, remember that we can recognise faces, on the basis of their intricate (unique) shapes, though we might struggle to describe the shapes of any of the face's features. The implication is that we can possess the concept of a shape without possessing the concepts of those shapes out of which it is composed. In which case, absent any alternative explanation, it cannot be an analytic truth that squares are composed of *straight lines*.

And once we have accepted all of the above, it becomes plausible to suppose that someone might grasp all of the relevant concepts, and might understand the sentence 'every square is a shape composed of four straight lines of equal length', and yet might refuse to assent to that statement. For example, they may have only recently acquired the concepts, and not yet noticed that squares are composed of straight lines. In which case, the proponent of the seeing-that view is deprived of the suggested explanation of the putative analyticity of such statements. There is no

⁵⁴ Though if the explanation that I offered, in the previous chapter, of our capacity to perceptually experience determinable properties is correct then she will not be able to do this either.

explanation of how, or why, one should be disposed to *assent* to such statements simply in virtue of *understanding* them. And since knowledge presupposes assent, this means that there is no explanation of how, or why, one might come to *know* such statements simply in virtue of *understanding* them. So the advocate of the seeing-that view is no closer to providing us with an explanation of where the relevant knowledge originates from. In the final section of this chapter we shall consider the only remaining explanation available to the proponent of the seeing-that view.

5.3. Innate Beliefs

The only remaining explanation available to the advocate of the seeing-that view seems to be to posit a class of innate beliefs pertaining to the various spatial composition relations. This will enable them to say that the relevant truths are a priori because, if the beliefs are innate, then one may hold the relevant beliefs even if one has not perceptually experienced any spatial composition relations. Of course, if those beliefs are to count as knowledge then the proponent of the seeing-that view will have to show that innate beliefs, of this kind, might be *justified* (or, rather, that the subject might be justified in holding such beliefs). And, as I shall mention, this further demand might pose a problem for them. But, as we shall see, the biggest problem with this explanation is that it forces us into a fresh confrontation with the phenomenological observations made in chapter 4.

On the assumption that a person cannot believe some proposition - that all bachelors are unmarried men, for example - without possessing the corresponding concepts - 'bachelors', 'marriage', and 'men', for example - the present proposal (innate beliefs), like the previous one (analytic truths), implies that the concepts of spatial composition and of various types of shape property (e.g. squares and straight lines) are innate⁵⁵. But, again, we don't just have to explain our capacity to see-*that* squares

⁵⁵ Crane (2009) distinguishes two notions of 'conceptual content'. On one, a content is conceptual if it is composed of concepts. On another, a content is conceptual just if one must possess the relevant concepts in order to hold the relevant belief. The motivation for the second notion is that some people think that belief contents are sets of possible worlds, and so are not composed of concepts. But Crane does not doubt that all will at least agree that one must possess the relevant concepts in order to hold

are composed of straight lines, we also have to explain our capacity to see-*that* a triangle is composed out of straight lines, and to see-*that* a circle is not. Indeed, the same follows for any pair of types of shapes that we can see (even those complicated shapes for which we have no name). The implication is that the proponent of the seeing-that only view must posit innate concepts for every type of shape that we can see (even those complicated shapes for which we have no natural language term). And, of course, the view also implies that all of these concepts must come already programmed into a massive range of beliefs, at birth. As such, the view is far from parsimonious.

This view, unlike the previous one, does not implicate any *constitutive* inter-conceptual relations, and it does not imply that a subject must know, or be in a position to know (through mere reflection), any particular definition, in order to count as understanding a word (and as grasping the corresponding concept). So the argument in the previous section has no traction against this view. However, the present view does imply that ordinary subjects (who are born with the relevant innate beliefs) will be in a position to offer such (canonical) definitions.

One very general epistemological worry with the view that the relevant beliefs are innate (the IB view) is that it is unclear whether the innate beliefs can be considered as justified. If they cannot then they may be incapable of conferring warrant upon our beliefs regarding the particular squares and straight lines that we see. This objection ought to concern those who hold an internalist and foundationalist theory of justification. If one is an internalist about justification then one will expect subjects to have access to that which justifies their beliefs. If one is a foundationalist about justification then one will expect justification to bottom-out with perceptual experience (and, perhaps, analytic truths). But if IB is the whole story about our knowledge of spatial composition relations then subjects will not have access to any perceptual beliefs that, however indirectly, might serve to justify those beliefs (and to constitute them as *knowledge*). So, anyone who is an internalist and a foundationalist about justification cannot endorse the present proposal.

the relevant belief, so, on that basis, Crane's second notion of conceptual content accommodates the intuition that belief contents are, in some sense, conceptual - even if they're sets of possible worlds.

Provided that one rejects either internalism or foundationalism about justification, and one is prepared to say that we have an innate concept for every type of shape that we can see (even those complicated shapes for which we have no name), and that all of these concepts come already utilised in a massive range of beliefs, at birth, then I think that this view can explain our knowledge of spatial composition relations. However, this view forces us into a fresh confrontation with the issues raised in chapter 4.

Recall, the proponent of the seeing-*that* only view does not claim that seeing-*that* the square is composed of straight lines in any way influences the phenomenal character of our perceptual experience. The issue here is not whether phenomenal-composition is the product of cognitive penetration, for example (I made the case against *that* view in section 3 of chapter 4); the issue is whether there is such a thing as phenomenal-composition at all. The proponent of the seeing-*that* only view denies that there is. So, on the seeing-*that* only view there is no reason to think that the phenomenal character of the perceptual experience of some subject who lacks the relevant innate beliefs will be in any way different from our own.

Suppose that someone did not have the relevant innate belief. After all, for any innate belief it is possible that a subject should be born without it (perhaps they lacked the innate belief gene!). Still, given the above, this subject would have perceptual experience exactly like our own. This means that we can turn to our own phenomenology in order to determine how plausible this view is.

Given that this subject's experience is just like our own, couldn't we make her understand that the square is composed of straight lines simply by directing her to attend, in the right way, to what she sees? The answer is surely: 'yes!' The reason being that, as a matter of fact, our visual experience does not seem to be silent on whether the square is composed out of the straight lines, or instead composed out of the colour, that it shares its location with. Visual experience makes manifest that it is the straight lines, and not the colour, that the square is composed of. This was precisely what we discovered in chapter 4 by way of a careful reflection on the phenomenology of our visual experience, and our modal intuitions regarding

different types of visual experiences/visualisation experiences (i.e. squares without straight lines and football stadiums full of faces replete with details).

The composition relation manifests itself through the phenomenal unity of the phenomenal-square and the phenomenal-straight-lines. As I have said, the unity of the phenomenal properties can make the diversity of elements easy to miss, but once we have a clearer view of those elements – the phenomenal-square, the phenomenal-straight-lines, and the relation between these -, it becomes very difficult to countenance the idea that we are not *visually aware* of an intrinsic connection between squares and straight lines.

Once again, the advocate of the seeing-*that* only view faces the set of choices that we observed in the previous chapter. They may deny that there are similarity relations among the phenomenal-shapes that reflect any similarity relations among the shapes simpliciter. Alternatively they may accept that there are, but deny that by virtue of these we perceptually experience the similarity relations among the shapes. It was noted in the previous chapter that if the proponent of the seeing-*that* only view opts for the latter then the similarity relations among phenomenal properties is epistemically redundant. But it was also noted that, since these similarity relations among phenomenal-shapes are introspectively accessible, the subject *could* make use of them.

The above example of the subject born without the innate composition-related beliefs helps to make this last point vivid. There do seem to be similarity relations among the phenomenal-shapes that reflect similarity relations among the shapes simpliciter. And even if the similarity relations among the phenomenal-shapes have thus far never caused some subject to make any similarity judgements about shapes, or to purposefully manipulate shapes, there seems to be no reason why they could not. Surely we can direct her attention to these similarity relations so that she might utilise them. As a result, it now just seems dogmatic to go on denying that this subject is perceptually experiencing the composition relation.

5.4 Summary

The conclusion of this chapter is that the advocate of the seeing-that view is faced with a dearth of plausible explanations for our knowledge of spatial composition relations. We have established that they are on extremely shaky ground if they wish to explain our knowledge of spatial composition relations in terms of analytic truths. This for two reasons:

1. As Williamson has shown, the very notion of 'analytic truth' looks like it may be misguided. This notion requires that the relevant truths exhibit understanding-assent links (not to mention understanding-justification links), which it seems that no truths in fact exhibit.

And even if there are such things analytic truths, (and, still, setting aside the issue of justification) the proponent of the seeing-that account needs to offer us an explanation of *how* the relevant understanding-assent links hold. The only candidate explanation of these involves appeal to the classical model of concepts. But when we consider the application of this model of concepts to the particular concepts that we are interested in – spatial concepts – we get the second reason.

2. If the classical model of concepts is correct then it seems that our spatial concepts must be observational concepts. In which case, possession of any of the relevant concepts does not implicate possession of any of the others, and so does not implicate any kind of (tacit) assent to propositions involving those latter concepts. This means that the model is powerless to explain how understanding-assent links might hold for statements about spatial composition relations, and so powerless to explain how such statements might be analytic.

The only remaining option available to the proponent of the seeing-that view is to say that the relevant beliefs are innate. This is likely to strike many as, intuitively, highly implausible. For one thing, the view is far from parsimonious, for it implies that a) all the concepts involved in the innate beliefs are innate, and b) those concepts

are innately programmed into all of the relevant beliefs. As we saw, there are also some theoretical issues around the question of whether such beliefs might count as *justified* (or, rather, whether the subject might count as *justified* in holding them), and, therefore, whether the beliefs might count as *knowledge*. But the principle objection to the innate beliefs proposal is that it forces us into a fresh confrontation with the phenomenological observations made in chapter 4, which observations tell against *any* seeing-that only account.

Chapter 6 - Knowledge of the Metaphysical Necessity of Spatial Composition.

In chapter 4 I sought to describe the intimate relationship between phenomenal-squares and phenomenal-straight-lines, which is manifest in perceptual experience. I described this feature of our perceptual experience as a phenomenal unity, and I suggested that it might serve to explain why we think of spatial composition relations as *metaphysically* necessary. So the case that I made there was that the phenomenology seems consistent with the view that we perceptually experience these relations *as* metaphysically necessary.

Mirroring the argument for perceptual experience of spatial composition, as it was presented in chapter 4, I could just say that, given this phenomenological fact, it would be dogmatic to go on denying that we perceptually experience the *metaphysical necessity* of composition relations. To deny that we perceptually experience the *metaphysical necessity* of composition relations would render the necessity of phenomenal-composition epistemically redundant. But since we are introspectively aware of the necessity of phenomenal-composition, what is to stop us from using the necessity of phenomenal-composition to inform us of the necessity of spatial composition? My opponent would have to say that there would be something epistemically deviant about doing so. But to say this would be dogmatic in the extreme.

I will not rest my case here. One reason is that I take it that the claim that I am now arguing in favour of will strike many as just that bit more contentious than the claim that we perceptually experience spatial composition⁵⁶. Another reason is that I think

⁵⁶ Moreover, even in the case of perceptual experience of spatial composition, which I take to be less contentious, I didn't ultimately rest my case with the phenomenological argument, for in the previous chapter I offered an epistemological argument for perceptual experience of spatial composition. So the epistemological argument that I am about to present for perceptual experience of modality exactly parallels that which I presented in the previous chapter, in favour of perceptual experience of spatial composition.

that the phenomenal-unity of phenomenal-squares and phenomenal-straight lines is just that bit subtler than are the similarity relations among phenomenal-shape-properties, which I appealed to in chapter 4. Even though the similarity relations between phenomenal-squares and phenomenal-triangles is *due* to their both exhibiting this phenomenal-unity with their phenomenally-composing phenomenal-straight-lines, the phenomenal-unity *itself* is somehow much harder to attend to (this is precisely why I took the indirect route of using the similarity relations in order to argue in favour of perceptual experience of spatial composition, rather than simply appealing directly to the phenomenal-unity). The reason for this is that, as I said in chapter 4, it is in the nature of unities that they resist attempts at articulation, for articulation implies division.

The obvious explanation for the advocate of the seeing-that only view to offer of our knowledge of the *necessity* of spatial composition relations is, once again, that the composition relations are analytic. In which case, they would be free to say that we know that the relations are necessary in virtue of our recognition of their analyticity. However, this route is not available to the proponent of the seeing-that only view, for it was refuted in the previous chapter.

In section 1 I will argue that there is at least a *prima facie* problem in accounting for our knowledge of the metaphysical necessity of composition relations. In section 2 I will argue that if we assume that we do perceptually experience the metaphysical necessity of composition relations (as our perceptual phenomenology seems to suggest) then we can understand our *knowledge* of its metaphysical necessity. I will also begin the process of constructing a positive proposal about what this involves. This positive proposal will involve appeal to an acquaintance relation. We shall consider alternatives to the acquaintance relation in chapters 7 and 8, and then, in chapter 9, we shall complete the task of constructing the positive proposal, with the use of the acquaintance relation. In section 3 of this chapter I respond to an objection to the view advanced in section 2.

6.1. The Source of our Knowledge of the *Metaphysical Necessity* of Spatial Composition Relations.

6.1.1 Induction

One option available to the advocate of the seeing-that only view is to say that we arrive at the belief that squares are *necessarily* composed of straight lines because every time anyone has ever seen a square, it has been composed of 4 straight lines. It is because one *never* sees a square that is not composed of 4 straight lines that we believe that squares are *necessarily* composed of straight lines. Perhaps our knowledge of the asymmetry of the composition relation can also be explained by the fact that we often see straight lines without squares (for sometimes they form triangles), but we never see squares without straight lines. However, there are two fatal problems with the idea that this knowledge is merely an empirical generalisation.

The first problem is similar to problems encountered when considering, in the previous chapter, the idea that statements about spatial composition relations might be analytic truths. It is unclear whether we could construct inductive hypotheses regarding complicated shapes for which we have no (non-demonstrative) concepts. And even if this is possible, the suggestion is undermined by our capacity to see-*that* unique shapes, which we have never before encountered, are composed of their parts – individual face shapes are an apt example.

The second, and more serious of the problems, is that we tend to think that the necessity that we believe to characterise spatial composition relations is of an altogether different kind to that which is established by induction. That is, we do not think of the necessity of the spatial composition relation as being of the same kind as other empirically established generalisations. Rather, we tend to think that spatial composition relations are characterised by *metaphysical* necessity, rather than just empirical/nomological necessity. But it is empirical necessities, not metaphysical necessities, which are established by induction.

In section 3, of chapter 4, I said:

‘Visual experience is not only sensitive to the existence of the composition relation, but it is also sensitive to the profound intimacy of that relation’

The intimacy that we are sensitive to, I am arguing, is metaphysical necessitation. If this is correct then clearly repeated experiences of squares are not *necessary* for the acquisition of the knowledge that squares are necessarily composed of straight lines. This necessity will present itself in a single visual experience of a square. I believe that I have already made a strong case for this claim in chapter 4. But I now wish to argue that if we assume that we do not perceptually experience the necessity of spatial composition then repeated experiences of squares will not even be *sufficient* for the acquisition of the knowledge that squares are necessarily composed of straight lines. At least, it will not be sufficient for knowledge of the kind of necessity that we take to be involved in spatial composition – that is, metaphysical necessity.

Most philosophers assume that the causal laws governing causal relations are merely empirically necessary, so there are possible worlds in which the causal laws fail and the same types of events can be otherwise causally related⁵⁷. But in the case of spatial composition we think that the necessity of the relation is metaphysical, so there is no possible world in which the same spatial properties can be otherwise compositionally related to one another (or not compositionally related to one another at all). It seems fair to say that, if a subject cannot grasp the metaphysical necessity of the composition relation on the basis of looking at one example of it, it is dubious that more examples will help.

What is to explain this difference in how we conceive of the necessity of causal relations and composition relations? I claim that the difference may be explained by the fact that in the case of composition, unlike mere natural laws, we see the relation! I say ‘may’, because some people think that we can perceive causal relations⁵⁸, and I don’t here wish to take a stand on the question of whether we can perceive such relations. But if we do perceive causal relations then I suggest that the difference

⁵⁷ C.B. Martin (2007), and Heil (2003) are notable recent exceptions. But their views are very much against the prevailing current.

⁵⁸ See Siegel (2010).

inheres in the fact that we perceive composed shapes as arising from the *intrinsic nature* of the composing shapes. The phenomenal-square arises from the intrinsic nature of the phenomenal-located-straight-lines, and it is in virtue of this fact that we *see* that the square arises from the intrinsic nature of the 4 located straight lines. If we do perceptually experience causal relations then we do not perceptually experience them *like that*, and this explains why we feel compelled to make a distinction between metaphysical, and merely empirical (or nomological), necessity.

6.1.2 Innate Beliefs, Take Two.

Mirroring the discussion in chapter 5, there is logical space for the view that the belief that spatial composition relations are metaphysically necessary is innate. As has been mentioned, there are general epistemological worries about the justification of innate beliefs. But this view raises a special problem. The view requires that the concept 'metaphysically necessary' is innate, and this seems implausible. What teleological pressure could possibly account for the selection and preservation of this concept? And what teleological pressure could possibly account for the selection and preservation of the innate belief that spatial composition relations are metaphysically necessary?

If there is an account to be offered of how such a belief might have been preserved it seems likely that it will have nothing to do with the truth of the belief. For example, it might be something along the lines of 'this belief made their possessors seem interesting to potential partners, thereby increasing reproduction' (not that this sounds particularly plausible either). When considering IB accounts of our knowledge of spatial composition relations, in chapter 5, I mentioned that such accounts are unamenable to those who hold foundationalist and internalist views about justification. But, if the above is on the right track, it looks as though even those who reject internalism about justification will be unable to avail themselves of the IB account of our knowledge of the necessity of composition relations.

Externalists tend to think that a belief is justified (or constitutes knowledge) in virtue of some kind a link between our possession of the belief and the truth of the belief.

This link takes different forms in different authors – for Goldman (1967) it's causal, for Nozick (1981) it's some kind of counter-factual dependence relation. But if the above is on the right track then it looks as though, regardless of the variant of externalism, there will not be a link of the right kind between our possession of the belief and the truth of the belief.

6.1.3. Conceivability

Chalmers (2002) and Yablo (1993) have both offered accounts on which conceivability/imagination may provide us with knowledge of modality. Chalmers and Yablo offer very similar accounts, but Chalmers' account introduces some distinctions that Yablo's account abstracts away from. It will be worth considering the additional distinctions that Chalmers makes, so I will focus on his version of the view. Chalmers (2002) attempts to demonstrate not only that we have knowledge of modal truths, but that we *a priori* knowledge of modal truths. He distinguishes two types of conceivability.

Primary conceivability:

'Primary conceivability is always an a priori matter. We consider specific ways the world might be, in such a way that the true character of the actual world is irrelevant. In doing so, empirical knowledge can be suspended, and only a priori reasoning is required' (2002: 158).

And secondary conceivability:

'Unlike primary conceivability, secondary conceivability is often a posteriori. It is not secondarily conceivable that Hesperus is not Phosphorus, but one could not know that a priori. To know this, one needs the empirical information that Hesperus is actually Phosphorus. This a posteriority is grounded in the fact that the application of our words to subjunctive counterfactual situations often depends on their reference in the actual world, and the latter cannot usually be known a priori'. (2002: 159)

According to Chalmers, primary conceivability is a guide to 'primary possibility', and

secondary conceivability is a guide to 'secondary possibility'. Officially, this distinction relies on Chalmers' view on which each belief or thought has two contents – one narrow and one wide. The idea is that the subject may have a priori knowledge of the narrow contents of her beliefs or thoughts, and so a priori knowledge of any implications of the narrow contents of her beliefs or thoughts. Chalmers defines narrow content, which he calls 'primary intensions', as follows:

'The primary intension of S is true at W if the material conditional 'if W is actual, then S' is a priori'.

However, this has some surprising implications:

The primary intension of some terms can vary between speakers. For example, Leverrier might use 'Neptune' to pick out whatever causes certain orbital perturbations within a world, whereas a friend might use it to pick out (roughly) whatever Leverrier refers to with the name, irrespective of any perturbing role' (2002: 167).

So the primary intension associated by one speaker with some sentence, S, may be different to that associated by another speaker with that sentence, S, and hence whether or not S is primarily conceivable may well vary between speakers. For Leverrier's friend it might be *primarily* possible that Neptune does not cause orbital perturbations of the relevant kind, whereas for Leverrier it will not be so.

This seems to be precisely the kind of situation that, as we saw in the previous chapter, Williamson warns us against. Recall Peter and Stephen, who deny that 'every vixen is a vixen'. Talking of the possibility that Peter and Stephen might associate idiosyncratic thoughts with that sentence, Williamson says:

'If Peter and Stephen associate (the sentence) with different thoughts from ours, should we not understand them better by translating their idiolects non-homophonically into ours? Presumably we should seek sentences other than (that sentence) that we associate with the very thoughts that they associate with (that sentence), or at least sentences we associate with thoughts more similar to the thoughts they associate with (that sentence), and translate the

dissent from (that sentence) in their mouths as dissent from those other sentences in our mouths' (Williamson, 2009: 114-115)

And he goes on to say:

'To insist on applying such a non-homophonic translation scheme to them in the face of their protests would be to treat them less than fully seriously as human beings, like patients in need of old-fashioned psychiatric treatment, whose words are merely symptoms. The claim that Peter and Stephen associate (that sentence) with different thoughts from ours repackages our disagreement with them in a way that makes it sound less threatening than it really is.' (Williamson, 2009: 115)

On the other hand, perhaps what is really objectionable in a view like Chalmers' is that the putative narrow content is actually *inexpressible*⁵⁹. In which case, the kinds of a priori modal truths that Chalmers is attempting to describe are really inexpressible too. This Chalmers might be happy to accept, if it means that we can hold on to some kind of a priori access to certain modal truths (albeit inexpressible ones). But the requirement that narrow contents do have interesting a priori implications places some further constraints on his notion of narrow content.

For example, when it comes to visual experience, Chalmers' (2006) narrow contents are all very similar to one another: the property that normally causes experiences like *this*, in this subject⁶⁰. Narrow contents like *that* would not have the kinds of implications that Chalmers is talking about. The only thing that could be inferred, a priori, on the basis of narrow contents like that, is that the relevant entities, if they exist, must be causally efficacious properties. For all that the narrow content implies, absolutely anything else (including all manner of modal statements) could be true of those properties. So, clearly, on Chalmers' view the narrow contents of

⁵⁹ See Farkas (2008).

⁶⁰ This is Chalmers' (2006) construal of the narrow content of a *colour* experience. See Thompson (2010) for the same construal of narrow content applied to visual experience of *shapes*.

beliefs/thoughts must somehow have more substantive a priori implications in some way built into them.

For our purposes we need only know if Chalmers can explain our knowledge that, necessarily, every square is composed of 4 straight lines of equal length. But, so far, he has provided us with no more resources for explaining this than we had when, in the previous chapter, we were considering the possibility that such knowledge might be analytic. Chalmers makes another distinction, which might furnish us with some further resources. This is the distinction between negative and positive conceivability.

Negative conceivability:

‘We can say that S is prima facie negatively conceivable for a subject when that subject, after consideration, cannot rule out S on a priori grounds.’ (2002: 149).

Positive conceivability:

‘Positive notions of conceivability require that one can form some kind of positive conception of a situation in which S is the case. One can place the varieties of positive conceivability under the broad rubric of imagination: to positively conceive of a situation is to imagine (in some sense) a specific configuration of objects and properties’. (2002: 150)

It is positive conceivability that Chalmers believes to be the stronger guide to metaphysical possibility. But it’s far from clear why our capacity to imagine a specific configuration of objects and properties should give any more credence to the idea that such a configuration should be metaphysically possible. And in the present case we are looking not for credence for the claim that something is metaphysically possible, but for credence for the claim that something is metaphysically impossible. Can our failure to positively conceive of a square not composed of 4 straight lines afford us with warrant for the claim that such is not metaphysically possible?

This would seem to get matters the wrong way around. In the case that we considered in chapter 4, in which we contrasted visualising a football stadium full of

faces, replete with details, all at once, and visualising a square without straight lines, we concluded that while it seems to be impossible to visualise either, we have different intuitions about why this is so. In the former case our failure seems to be symptomatic of a contingent shortcoming in our visualising capacities. In the latter case we are inclined to instead say that the task is metaphysically impossible. Chalmers (and Yablo) are unable to cast any light on this asymmetry.

What if we focus instead on *secondary* conceivability? Chalmers believes that what secondary conceivability adds to primary conceivability is knowledge of which worldly properties, kinds, and individuals her words/concepts in fact apply to. For example, it adds the knowledge that 'water' refers to the same stuff as 'H₂O', and 'Hesperus' refers to the same object as 'Phosphorus', and 'square' refers to the shape composed of 4 straight lines of equal length. This, Chalmers believes, provides additional constraints on our exercise of what he calls 'modal imagining'. The idea is that this knowledge limits what is secondarily conceivable so that one finds one's self unable to 'modally imagine' that Hesperus is distinct from Phosphorus, or that water is distinct from H₂O, or that squares are composed otherwise than by 4 straight lines of equal length.

Roca-Royes observes 'There should be no doubt that essential properties are what secondary intensions consist of. H₂O is assumed to be the essence of water and it is also assumed to be, by two-dimensionalists, the secondary intension of water' (2011: 43). The implication, Roca-Royes concedes, is that '*H₂O is not water*' is secondarily conceptually contradictory. But in that case the knowledge added by *secondary* conceiving already assumes essentialist knowledge. In order to know that 'Hesperus' refers to the same object as 'Phosphorus', we must know that Hesperus has the same essence as Phosphorus. In order to know that 'water' refers to the same stuff as 'H₂O', we must know that water has the same essence as H₂O. In order to know that 'square' refers to the shape composed of 4 straight lines of equal length, we must know that squares have the same essence as shapes composed of 4 straight lines of equal length. So, if our knowledge that squares are *necessarily* composed of 4 straight lines comes from secondary conceivability then this knowledge is grounded in *essentialist* knowledge.

As Roca-Royes puts it:

Chalmers 'intends us to believe that the SIC (secondary ideal conceiver) knows about secondary intensions merely by knowing all non-modal facts about the actual world. This cannot be the case. For surely not all the non-modal facts that the SIC knows enter secondary intentions - in particular, contingent facts must not and all essential ones must. So the SIC must know in addition which non-modal facts do and which ones do not enter secondary intentions. By primitively providing the SIC with knowledge of secondary intensions, therefore, Chalmers is providing the SIC with non-elucidated essentialist knowledge' (2011: 43)

Lowe (2012) is reluctant to assent to the identification of water with H₂O, precisely because, in his view, this would have the implication that water and H₂O have the same essence.

'Any sample of pure water contains OH⁻ and H₃O⁺ ions as well as H₂O molecules, which explains why even pure water conducts electricity, albeit only very weakly; so, if 'Water is H₂O' is understood as asserting water is identical with, or is wholly composed by, H₂O molecules, then it is simply false. (Lowe, 2012: 920)

So it is certainly possible to dispute whether or not water is identical to H₂O, even with abundant knowledge of the relevant non-modal truths. The implication of Roca-Royes's and Lowe's observations are that even if we, somewhat unrealistically, idealise the subject's non-modal (and non-essential) knowledge (as Chalmers explicitly does), this alone will not suffice to generate *any* modal knowledge. The principle difference between the way in which our 'modal imagining' is meant to be constrained in the primary and secondary modes is that in the secondary mode, unlike the primary mode, we hold fixed identity/constitutive/compositional facts⁶¹. But the crucial

⁶¹ Another potential constraint that's been discussed in the literature is origins. Kripke (1980) famously claimed that a person's origins are essential to them. If this is correct then when we are secondarily conceiving we ought to keep facts about origins held fixed. But how do we know if Kripke is correct? As Roca-Royes (2011) points out, without circularity we cannot appeal to the results of secondary conceiving to answer this question.

question, to which neither Yablo nor Chalmers offer any answer, is how do we know these identity/constitutive/compositional facts?

Chalmers and Yablo could deny – what Roca-Royes and Lowe both assume - that the words ‘constitute’, ‘identical’, and ‘compose’, are words that have any immediate implications concerning essence or modality. In which case it would not be objectionable that Chalmers appeals to knowledge of such facts as a constraint on our ‘modal imagining’. But then these facts don’t seem suited to playing the guiding/constraining role that’s needed of them. If constitutive facts are not essentially facts about essence or modality, how does one know to keep constitutive facts held fixed? And what justifies this practice? So this line looks untenable.

There is a final issue for the imagination-based accounts, which serves to introduce the next proposal to be considered. Williamson (2007) argues that imagination-based accounts seem to implicate us with a *sui generis* capacity for generating a kind of knowledge that has no survival value. This is quite vivid when Chalmers refers to the practice that he describes as ‘modal imagining’. Williamson puts forward his own account, which is intended to avoid this objection. We shall turn our attention to Williamson’s account in the following section.

6.1.4 Knowledge of Counter-Factuals.

Williamson (2007) has propounded his own account of our modal knowledge. I will argue that his account, as it stands, is generally unviable as an account of our knowledge of metaphysical (as opposed to merely empirical) necessity. However, if supplemented with *some* perception-based knowledge of some essential truths then Williamson’s model may provide a route from the limited stock of modal truths afforded by perceptual experience (including, at least, truths pertaining to the necessity of spatial composition relations) to a much richer stock of modal truths. So Williamson’s account does not provide an alternative to perceptual experience of the necessity of spatial composition relations. Indeed, the gaps in Williamson’s account actually point towards the existence of perceptual experience of the necessity of spatial composition relations.

Talking of our knowledge of modal truths (a distinctively philosophical type of

knowledge), Williamson (2007) says:

‘We should expect the cognitive capacities used in philosophy to be cases of general cognitive capacities used in ordinary life, perhaps trained, developed, and systematically applied in various special ways, just as the cognitive powers that we use in mathematics and the natural sciences are rooted in more primitive cognitive powers to perceive, imagine, correlate, reason, discuss... In particular, a plausible non-skeptical epistemology of metaphysical modality should subsume our capacity to discriminate metaphysical possibilities from metaphysical impossibilities under more general cognitive capacities used in ordinary life’ (2007: 136).

Williamson believes that our knowledge of modal truths has its source in our capacity for counterfactual reasoning. Williamson suggests that it is uncontroversial that we have the capacity to make real-time predictions about our environments. As one watches a boulder roll down a rock face one instantly and effortlessly forms expectations about where the boulder will end up. This is what makes it possible for us to take effective evasive action, and to avoid being squashed by the boulder. Williamson calls this capacity ‘simulation’. It is Williamson’s contention that, when evaluating counterfactuals, we use this same capacity, but we use it ‘offline’. This means that the capacity is not constrained to make predictions only about actual events. Instead, when the capacity is used ‘offline’, the subject is free to set alternative (counterfactual) initial conditions, and to make predictions about what would happen if those initial conditions were satisfied.

‘Our overall capacity for somewhat reliable thought about counterfactual possibilities is hardly surprising, for we cannot know in advance exactly which possibilities are or will be actual. We need to make contingency plans. In practice, the only way for us to be cognitively equipped to deal with the actual is by being cognitively equipped to deal with a wide variety of contingencies, most of them counterfactual’. (2007: 137)

Williamson believes that our knowledge of metaphysical modality can also be explained by the offline exercise of simulation. If this account is viable then it has the merit of providing a naturalistic and unified account of a range of epistemic

capacities:

'Far from being *sui generis*, the capacity to handle metaphysical modality is an "accidental" byproduct of the mechanisms that provide our capacity for ordinary thinking about the natural world, which involves counterfactual thinking, skeptics about metaphysical modality cannot excise it from our conceptual scheme without loss to ordinary thought about the natural world, for the former is implicit in the latter.' (2007: 162)

It is central to Williamson's account that much of what constrains our simulations when engaging in counterfactual reasoning is a kind of folk theory of physics. Crucially, the 'folk physics' is not to be thought of as a body of propositional knowledge. The reason for this is that, as Williamson admits, the 'theory', though locally reliable, would almost certainly be false. If the folk-physics were regarded as propositional then, given that it would be false, any potential 'knowledge' of counterfactuals obtained from it would be thereby barred from actually counting as knowledge.

We should understand the 'folk theory' as more like a heuristic than a genuine theory. The heuristic is designed to enable us to deal with the macroscopic environment on earth. The folk physics will likely be 'inferentially insulated', in that it is unlikely that we can use it for tasks other than simulation, and it may not be susceptible to change in light of learning – for example, it won't change in response to a physics class, even if the subject's propositional knowledge regarding physics is much improved by the class. In addition to some kind of non-propositional folk physics, simulations will also draw on any, and all, propositional knowledge that the subject has. For example, the knowledge that water is identical with H₂O.

We are now in a position to understand Williamson's proposal with regard to knowledge of metaphysical modality. He suggests that we establish S as a metaphysical impossibility if simulations in which we counterfactually suppose S robustly yield a contradiction, e.g. S and not S.

The problem with Williamson's proposal is that the application of simulation to knowledge of metaphysical modality is dubious. Though Williamson is correct that

evolution does occasionally throw up “accidental” byproducts, on Williamson’s view it does become a mere article of faith that capacities, which evolved for certain practical purposes, remain reliable when used in a context that has no practical import at all.

As Lowe puts it⁶²:

‘The root trouble is that Williamson’s theory is based on a purely formal explication of modal propositions in terms of counterfactual ones, whereas the facility with counterfactuals that he appeals to in his account of how we can acquire modal knowledge draws only on our ability to handle causal counterfactuals, which are just not relevant where the metaphysical modalities are concerned’ (2012: 933).

In a similar vein, Roca-Royes (2011) points out that essentially the same objection that she raised to Chalmers’ conceivability account can be applied also to Williamson’s proposal.

‘According to Williamson’s epistemology of counterfactuals, when evaluating whether p (square, arrow, contradiction), some items of our background knowledge must be imagined away (thereby ceasing to be exploitable), and some others must be held fixed. Potentially anything that we hold fixed, and only what we hold fixed, can be exploited in order to, in conjunction with p - the counterfactual supposition - , arrive at a contradiction. Intuitively, therefore, we should expect that, to obtain the extensionally right results, this must be the case:

(HF) Whenever a q , from our background knowledge, is inconsistent with p : q is held fixed iff it is a constitutive fact.

The reason is as follows. If, in counterfactual evaluation, we held fixed, no matter what, that *my left arm is not broken*, a contradiction would counterfactually follow from *my left arm is broken*, from which we

⁶² For various arguments against Williamson’s proposal, which engage with the detail of Williamson’s counterfactual logic, see Lowe (2012).

would erroneously conclude the corresponding impossibility'. (2011: 37-38)

Simulation of a counterfactual situation in which water is not constituted by H₂O will generate a contradiction, on Williamson's view, on the assumption that constitutive facts are held fixed between possible worlds (and likewise with identity and composition facts - e.g. Hesperus and Phosphorus, squares and straight lines). Perhaps constitutive facts *are* embedded in the folk-physics that constrains our simulations. In which case, since Williamson does not think that the folk-physics is propositional, simulations might be sensitive to modal facts, without yet assuming that the subject already has any essentialist (propositional) knowledge.

Certainly, if the assumptions which are embedded into the folk-physics really do enable us to track constitutive facts, and thereby essences, then they would provide knowledge of what is possible in worlds in which the laws of physics are different from the actual world. But there is no reason to assume that they really will track constitutive facts. Presumably whatever constrains our simulations when we engage in counterfactual reasoning does so because it helps to produce the right results in the kinds of cases for which our capacity for counterfactual reasoning evolved. Williamson explicitly concedes that the folk-physics will likely be inaccurate. This is precisely why he maintains that the folk-physics is non-propositional - for were the folk physics propositional then its falsity would preclude it from being a source of modal *knowledge*. Williamson believes that the folk physics can generate the right results, in practical circumstances, despite its strict falsity. But in that case, why should we believe that those assumptions that constrain our simulations when we engage in counterfactual reasoning really do track constitutive facts, and thereby essences?

If the assumptions that constrain our simulations when we engage in counterfactual reasoning do not track constitutive facts, and thereby essences, then there is no reason to think that counterfactual reasoning will produce the right results in situations for which the capacity did not evolve (e.g. enquiries into metaphysical modality). Certainly, we did not evolve under any pressure to *know* worlds in which the actual laws of physics are violated. We did not evolve under any pressure to deal

with such worlds in any way whatsoever. As Lowe (2012) points out, what Williamson requires is that we be in possession of a (reliable) folk-*metaphysics*, not just a (reliable) folk-*physics*. However, as Lowe also points out, a folk-*metaphysics* is exactly what Williamson denies that we have. And Williamson is quite right to deny that we have any such thing, since there is no evolutionary pressure that would explain our possession of a (reliable) folk metaphysics. But this leaves Williamson unable to explain our knowledge of modal truths.

Williamson is here attempting to account for the entire gamut of metaphysical modal knowledge. In contrast, I am arguing only that we perceptually experience a certain class of spatial relations as metaphysically necessary. It is plausible that the capacity for counterfactual thinking is, as Williamson claims, crucial to our capacity to know many modal truths. But what we have found is that, if it is to play such a role, it must first be supplemented with at least some modal/essential truths. Without some such truths the capacity for simulation is likely to be under-constrained, and unable to track metaphysical modal facts. So perhaps a role for Williamson's simulations is to *extend* our stock of modal knowledge. This leaves a role for perceptual experience in providing some initial modal truths.

Indeed, some of what Williamson says is congenial to this approach:

'There is no uniform epistemology of counterfactual conditionals. In particular, imaginative simulation is neither always necessary nor always sufficient for their evaluation, even when they can be evaluated. Nevertheless, it is the most distinctive cognitive feature of the process of evaluating them, because it is so much more useful for counterfactuals than for most non-counterfactual contents, whereas reasoning, perception and testimony are not generally more useful for counterfactuals than for non-counterfactual contents' (2007: 152)

The above passage leaves open the question as to how much work these different cognitive powers – most relevantly, perception – are doing in explaining our capacity for counterfactual thought. The phenomenological observations that we made in chapter 4 suggest that perceptual experience may well provide at least some modal knowledge as *input* to the simulation process – i.e. modal knowledge of the necessity of spatial composition relations. Williamson's model may then provide a route from

the limited stock of modal truths afforded by perceptual experience to a much richer stock of modal truths.

Roca-Royes (2011), Lowe (2012), and Fine (1994) all believe that knowledge of modal truths has its source in knowledge of *essence*. Moreover, Roca-Royes (2011) explicitly concedes that our capacity for imagination/counterfactual reasoning might provide us with knowledge of modal truths if that faculty were antecedently stocked with knowledge concerning *essences*. In the next section I will put forward a view on which visual experience of some spatial properties affords us with knowledge of the essences of those properties.

6.2. A Positive Proposal

6.2.1 Knowledge of Modality Originates in Knowledge of Essence

Lowe (2012) argues that knowledge of metaphysical modality has its source in knowledge of essence. He considers two views of what essences are:

1. The object (type) itself
2. A 'real definition' (a definition of the object).

Lowe favours real definitions. He describes a real definition as follows:

'A real definition of an entity, E, is to be understood as a proposition which tells us, in the most perspicuous fashion, what E is—or, more broadly, since we do not want to restrict ourselves solely to the essences of actually existing things, what E is or would be. This is perfectly in line with the original Aristotelian understanding of the notion of essence, for the Latin-based word 'essence' is just the standard translation of a phrase of Aristotle's which is more literally translated into English as 'the what it is to be' or 'the what it would be to be' (Lowe 2008a, p. 35)' (Lowe, 2012: 935).

So a real definition is a definition of an object (type), and this is to be distinguished from a 'verbal definition', which is a definition of a word. Lowe borrows this distinction from Fine:

‘It has been supposed that the notion of definition has application to both words and objects-that just as we may define a word, or say what it means, so we may define an object, or say what it is. The concept of essence has then taken to reside in the "real" or objectual cases of definition, as opposed to the "nominal" or verbal cases’. (Fine, 1994: 2)

Moreover:

‘There is more to the idea of real definition than is commonly conceded. For the activities of specifying the meaning of a word and of stating what an object is are essentially the same; and hence each of them has an equal right to be regarded as a form of definition’ (Fine, 1994: 14).

For Fine, both kinds of definitions involve identifying essences; it’s just that in the one case – verbal definition - it is the essence of a meaning (or, if words are given a thick individuation, the essence of a word), and in the other case – real definition - it is the essence of an object or property.

Fine’s account stands opposed to the presently dominant practice of attempting reduce essential truths to necessary truths:

‘The notion of essence which^[SEP] is of central importance to the metaphysics of identity is not to be understood in modal terms or even to be regarded as extensionally equivalent to a modal^[SEP] notion. The one notion is, if I am right, a highly refined version of the other; it is like a sieve which performs a similar function but with a much finer mesh’ (Fine, 1994: 3).

And Fine offers some examples of necessary, but inessential, truths:

‘Consider, then, Socrates and the set whose sole member is Socrates. It is then necessary, according to standard views within modal set theory, that Socrates belongs to singleton Socrates if he exists; for, necessarily, the singleton exists if Socrates exists and, necessarily, Socrates belongs to singleton Socrates if both Socrates and the singleton exist. It therefore follows according to the modal criterion that Socrates essentially belongs to singleton Socrates. But, intuitively, this is not so. It is no part of the essence of Socrates to belong to the

singleton. Strange as the literature on personal identity may be, it has never been suggested that in order to understand the nature of a person one must know to which sets he belongs. There is nothing in the nature of a person, if I may put it this way, which demands that he belongs to this or that set or which demands, given that the person exists, that there even be any sets.

It is not critical to the example that appeal be made to an abstract entity. Consider two objects whose natures are unconnected, say Socrates and the Eiffel Tower. Then it is necessary that Socrates and the Tower be distinct. But it is not essential to Socrates that he be distinct from the Tower; for there is nothing in his nature which connects him in any special way to it'. (1994: 4-5)

Fine adds that 'it lies in the nature of the singleton to have Socrates as a member even though it does not lie in the nature of Socrates to belong to the singleton' (1994: 5) but that 'no corresponding modal asymmetry can be made out' (1994: 5). Moreover, Fine identifies further problems for the programme of reducing essential truths to modal truths:

'Consider any necessary truth; it could be a particular mathematical truth, for example, or even the conjunction of all necessary truths. Then it is necessarily the case that this truth should hold if Socrates exists. But it is no part of Socrates' essence that there be infinitely many prime numbers or that the abstract world of numbers, sets, or what have you, be just as it is'. (1994: 5)

And he evinces that:

'It will part of the essence of any object that every other object has the essential properties that it has: it will be part of the essence of the Eiffel Tower for Socrates to be essentially a person with certain parents, let us say, or part of the essence of Socrates for the Eiffel Tower to be essentially spatio-temporally continuous. O happy metaphysician! For in discovering the nature of one thing, he thereby discovers the nature of all things' (1994: 6).

We begin to see Fine's positive proposal emerge from the following passage:

'What is it about the concept of necessity which makes it so inappropriate for

understanding the concept of essence? Certainly, there is a connection between the two concepts. For any essentialist attribution will give rise to a necessary truth; if certain objects are essentially related then it is necessarily true that the objects are so related (or necessarily true given that the objects exist). However, the resulting necessary truth is not necessary simpliciter. For it is true in virtue of the identity of the objects in question; the necessity has its source in those objects which are the subject of the underlying essentialist claim' (Fine, 1994: 8-9).

And he adds:

'Each object, or selection of objects, makes its own contribution to the totality of necessary truths; and one can hardly expect to determine from the totality itself what the different contributions were' (Fine, 1994: 9)

Lowe (2012) offers his own example of the difference between an essential truth and a merely (metaphysically) necessary truth, which is of particular interest to us, as it concerns geometrical properties:

'(E1) An ellipse is the locus of a point moving continuously in a plane in such a fashion that the sum of the distances between it and two other fixed points remains constant.

(E2) An ellipse is the closed curve of intersection between a cone and a plane cutting it at an oblique angle to its axis greater than that of the cone's side'.
(2012: 936)

According to Lowe, the first of these is an essential truth⁶³; but the second is a merely (metaphysically) necessary truth. One way of seeing this is by observing that:

'An ellipse can exist even in a purely two-dimensional space, but a cone can exist only in a space of at least three dimensions — hence it cannot be right to

⁶³ Incidentally, Lowe (2012) claims that relations between colours are a part of their essence. The relations that Lowe is talking about are those that we established in chapter 2 as colour composition relations. Johnston (1992) and Allen (2016) are also both in agreement that such relations are essential to the colours.

define an ellipse in terms of its relationship to a cone, since ellipses can exist perfectly well without cones' (2012: 937)

But, Lowe claims, we can see that the merely (metaphysically) necessary truth (E2) holds in virtue of the essential truth (Ei) (at least, it holds in virtue of (Ei) in concert with some further essential truth concerning cones). Of (E2) Lowe says:

'This metaphysically necessary truth holds in virtue of the essences of an ellipse and a cone, which are two quite distinct essences. It is because of what an ellipse is, and what a cone is, that this relationship necessarily holds between ellipses and cones. But it is not part of anything's essence that it holds' (2012: 939)

Lowe draws the following conclusion:

'Metaphysically necessary truth is a truth which is either an *essential* truth or else a truth that obtains *in virtue of* the essences of two or more distinct things. On this account, all metaphysical necessity (and by the same token all metaphysical possibility) is *grounded* in essence'. (2012: 939)

Neither Lowe nor Fine offers an account of how we come to know real definitions. Though Lowe does explicitly reject the idea that it involves any kind of 'acquaintance' relation. Part of the reason that Lowe denies that knowledge of essences involves acquaintance is because he does not think that essences are entities (which we might be acquainted with), but that they are instead real definitions. But we can accept this view and simply say that one grasps the real definition in virtue of being perceptually acquainted with the objects or properties to which they apply. This requires that, when one is acquainted with a property, at least some of that property's essential properties (in the case of a spatial property, this will include its spatial composition relations) feature in the experience's NR-content, and, moreover, that the necessity of those (essential) properties also feature in the experience's NR-content.

If we now look back at chapters 2 and 3, we can see that there is a rich story to be told about how our acquaintance with objects provides us with knowledge of the essence of an object in general. The word 'object' can be seen as a sortal, referring to a very

general – perhaps the most general – natural kind. In chapters 2 and 3 we saw how the identity of a particular object manifests itself in visual experience. We also saw how, on a Naïve Realist account, our perceptual experience of particular objects affords us with knowledge of the identity conditions of an object in general. We have now seen that an essence of some entity, E, is merely a real definition of E, where a ‘real definition’ of E is ‘what E is or would be’ (Lowe, 2012: 935). So the identity conditions of an object in general are (at least part of) the essence of an object in general. This means that the position propounded in chapters 2 and 3 offers us a detailed story about how perceptual experience, qua an acquaintance relation, is able to afford us with knowledge of (at least part of) the essence of an object in general.

When considering objections to her argument against existing accounts of modal knowledge, Roca-Royes muses over the following suggestion: ‘conceivability-based accounts do explain modal knowledge: they tell us that we obtain modal knowledge from essentialist knowledge’ (2011: 42), but, she replies: ‘I am identifying the following substantial explanatory deficit: on non-epistemic accounts, modal knowledge depends on a particular kind of knowledge, namely essentialist knowledge (whether modal or not). This knowledge is not elucidated and it *must* be’ (2011: 42). My account is that at least some such knowledge – that concerning the essences of certain spatial properties, and that concerning the essences of objects in general - is provided by perceptual experience, in virtue of an acquaintance relation.

We have been considering what the proponent of the seeing-that only view might say about our knowledge of the (metaphysical) necessity of spatial composition relations. This led us to discuss our knowledge of (metaphysical) necessity more generally. I have argued that conceivability/imagination-based accounts of such knowledge need supplementing with an explanation of where essentialist knowledge comes from. I wish to suggest that, at least in some cases (including our knowledge of the essences of spatial properties), the source of our essentialist knowledge lies in perceptual experience, and, more specifically, in *acquaintance* with the relevant properties.

6.2.2. – The Epistemic Role of Acquaintance

Though I have just proposed that we can explain our knowledge of some essential

truths by appealing to an acquaintance relation, nothing so far said favours this account over an account on which we perceive the necessity of spatial composition but on which perceptual experience is instead understood as the Intentionalist understands it.

Focusing on the necessity of spatial composition relations, I will consider the prospects for I-content accounts in chapters 7 and 8. Then, in chapter 9, I will elaborate the NR-content account, (which account I defined in chapter 1 as involving an acquaintance relation). However, there remains a question about whether an Intentionalist might not avail themselves of that very explanation. For, as I mentioned in the introduction, Tye is an Intentionalist, and he has recently (2009) endorsed the view that perceptual experience does involve an acquaintance relation. In which case, might not an Intentionalist, like Tye, espouse the view that our knowledge of the necessity of spatial composition relations has its source in an acquaintance relation?

In this section we shall give some consideration to what the nature of the acquaintance relation must be if it is to serve the suggested purpose. It will transpire that an acquaintance relation may only provide us with propositional knowledge if it serves to secure NR-content. As I defined the views in chapter 1, a view only counts as Intentionalist if it denies that perception involves *any* NR-content. Hybrid accounts, on which perception involves both I-content and NR-content⁶⁴, are to count as variants of Naïve Realism. The implication is that no Intentionalist is in a position to utilize the above proposal. This conclusion will set the scene for the subsequent discussion in chapters 7-9.

I mentioned in the introduction that Tye (2009) has recently endorsed the distinction between knowledge by acquaintance and knowledge by description. Now, Tye's acquaintance relation is crucially different from the Naïve Realist's because, for Tye, acquaintance with some property is realised when the subject undergoes an experience that *represents* that property. Whereas, for the Naïve Realist, the relation of acquaintance is primitive. But Tye does think that, through perceptual experience, we are *acquainted* with mind-independent properties.

⁶⁴ See Soteriou (2013).

Can Intentionalism, supplemented with an acquaintance relation, provide for our knowledge of essence, and, consequently, for our knowledge of the necessity of spatial composition? It seems not. Tye's knowledge by acquaintance is designed purely to accommodate his transparency thesis. There is no suggestion, from Tye, that knowledge by acquaintance can serve as the basis for *propositional* knowledge. Indeed, on Tye's account the acquaintance relation is merely something that supervenes on perceptual experience, and which explains its phenomenology. The perceptual experience proper is the representational state – that *in virtue of which* one is acquainted with the objects of perception –, and *that* (the representational state) is the basis of our perceptually-based propositional knowledge. More specifically, for Tye it is the I-content of the representational state that affords us with propositional knowledge.

In contrast, the Naïve Realist appeals to the acquaintance relation as a primitive, which is itself the basis for our propositional knowledge. For the Naïve Realist, all propositional knowledge acquired on the basis of perceptual experience is acquired on the basis of the acquaintance relation, for that is what perceptual experience is. So there is no problem about saying that we acquire knowledge of essence, in particular, on the basis of knowledge by acquaintance.

Tye could allow that the acquaintance relation is a source of propositional knowledge, but then he would have to concede that there is NR-content, in addition to I-content. At that point the only thing distinguishing Tye's view from a hybrid version of Naïve Realism⁶⁵ would be that Tye does not think that we are acquainted with particulars⁶⁶.

⁶⁵ See Soteriou (2013).

⁶⁶ Looking at the explanation that I offered, in chapter 1, of the Naïve Realist mechanism of content-fixation, it can seem as though the difference between Naïve Realism and the above development of Tye's view is greater than it really is. This is because I said that, for the Naïve Realist, what ultimately determines the content of a perceptual experience is the subject's environment. But on the above development of Tye's view it is still the I-content of an experience that determines which properties the subject is acquainted with (and therefore which properties feature in the experience's NR-content). But this is just a corollary of the fact that, on this development of Tye's view, it is still not particulars that the subject is acquainted with. For if it were, then those particulars would surely be the same particulars that fall in the subject's visual field, and it would surely be those particulars for that very reason (i.e. *because they fall in the subject's visual field*), so it would once again be the environment

For Tye, we are perceptually acquainted only with property types (universals?), so it would be property types (universals?) that featured in the NR-content.

If Tye is to accommodate perceptual experience of the necessity of composition relations, without making the above concessions, then he will have to show that this is something that can feature in the I-content of perceptual experience. We shall consider whether existing theories of I-content fixation are capable of delivering such contents in chapters 7 and 8. In the final section of the present chapter we shall consider an objection to the proposal that we acquire knowledge of the necessity of spatial composition relations from perceptual experience.

6.3. Objection – Revelation.

Revelation is a thesis that has been roundly rejected by philosophers of perception (Campbell: 2005, Allen: 2016, Shoemaker: 2006). Even Johnston (1992), who first propounded Revelation, as a common sense belief about colour, does not believe that it is strictly speaking true. In reading discussions of Revelation one can often sense an ambiguity in the notion of Revelation at issue. Sometimes it seems as though Revelation is a thesis regarding what it is that manifests itself in perceptual experience. But formal definitions of Revelation tend to construe it instead as a thesis about what propositional knowledge is afforded by perceptual experience, and arguments against Revelation invariably rely on this construal⁶⁷.

In agreement with this, Allen writes ‘Revelation is normally taken to involve ascribing propositional knowledge about the essential natures of the colours to subjects of experience: knowledge that the colours are essentially some way’ (2016, 131). Allen formulates a version of Revelation tied specifically to the colours:

‘REVELATION: the essential natures of the colours are fully revealed by visual experiences as of coloured things’ (131).

that determined the NR-content of the perceptual experience. The alternative way of getting particulars into Tye’s acquaintance relation, without giving the environment this content-fixing role, is to make the particulars mind-dependent (i.e. Sense Data Theory).

⁶⁷ See the chapter on Revelation in Allen’s (2016) for a number of arguments against Revelation, of this form.

Tye, who endorses the idea that, through perceptual experience, we are *acquainted* with the properties of mind independent objects, explicitly denies that this commits him to Revelation. He quotes a passage from Russell:

‘The particular shade of colour that I am seeing may have many things said about it—I may say that it is brown, that it is rather dark, and so on. But such statements, though they may make me know truths *about* the colour, do not make me know the colour itself any better than I did before: so far as concerns knowledge of the colour itself, as opposed to knowledge of truths about it, I know the colour perfectly and completely when I see it, and no further knowledge of it itself is even theoretically possible. (ibid., pp. 46–47)’ (2009: 97).

And of this passage Tye remarks:

‘In this passage Russell seems to be accepting what is often called “the thesis of revelation”—the thesis that colors have no hidden nature. There is, however, a weaker interpretation of the passage: that knowledge of a particular shade of brown via direct awareness of it is knowledge of a sort that cannot *itself* be improved or deepened by knowing truths about that shade of brown (which is not to say that one cannot add to such knowledge a further kind of knowledge of the given shade). In this way, knowledge by acquaintance of the color is complete and perfect... Consistent with these remarks, it could be held simply that knowledge of what that nature is, knowledge that the given shade of brown has a certain nature is not knowledge that one can glean by direct awareness. The latter is knowledge by description’ (2009:97)’.

In a similar vein, he later says:

‘Consider again the intense blue color of the computer monitor screen which I am staring. It seems to me that it is conceivable that the color which I am conscious in undergoing my present experience is not the color blue. Conceivably I am misclassifying the color. Conceivably I am wrong in believing

that *that* color is the color blue. But this, of course, is a failure to know with certainty a particular truth. It is not a failure to know a thing. And in failing to know with certainty the relevant truth, I do not fail to be conscious of the color I am experiencing. Thus, I do not thereby fail to know the entity I know. I still know it simply by being directly conscious of it. I know it by acquaintance' (2009: 189).

Unlike Tye, I do maintain that (propositional) knowledge of essence can be, and often is, afforded by perceptual experience, but this does not mean that such knowledge is invariably actually acquired on the basis of such perceptual experiences. Nor does it mean that when one does acquire knowledge of the essence of some property – say, a colour - through perceptual experience, then, as a result, that colour may 'have no hidden nature' (see the second passage quoted above). Something similar is implied by Allen's formulation of Revelation (concerning colour), above. According to Allen, Revelation (concerning colour) has it that 'the essential natures of the colours are *fully* revealed by visual experiences as of coloured things' (italics added)⁶⁸. But there is no reason for us to suppose that this is so. Finally, there is no reason for us to assume - as Tye implies in the third quoted passage that Revelation entails - that perceptual experience ever provides us with *certain* knowledge of essence. Mistakes, and therefore (reasonable) doubt, is always possible. My argument requires only that we do acquire some essentialist knowledge through perceptual experience, not that we acquire any *certain* essentialist knowledge through perceptual experience.

Let's consider some of the above points in more detail. If perceptual experience is capable of affording us with knowledge of essence, are we thereby entitled to assume that such experiences will *invariably* result in such essentialist knowledge? The Campbellian interpretation of Huang and Pashler's model of attention, which I presented in chapter 2, entails that the content of perceptual experiences frequently (usually even) outstrips the knowledge that we in fact acquire through those perceptual experiences. For, according to that interpretation, our visual experience involves an array of properties, stretching across the visual field, but we only acquire

⁶⁸ And that Revelation implies that essences are *fully* revealed by perceptual experience is crucial to all of Allen's (2016) arguments against Revelation.

knowledge of those properties that we go on to access.

Given that the content of perceptual experiences frequently (usually even) outstrips the knowledge that we in fact acquire through those perceptual experiences, we ought not to be surprised that, though perceptual experience makes knowledge of essence available to us, we don't invariably acquire such knowledge on the basis of those experiences. And it seems highly plausible that even when we do visually access some property there is no guarantee that we will acquire any, let alone all, essentialist knowledge concerning that property. For example, someone can perceptually experience the shape of a face (and, as I have argued, even simple shapes like squares) and yet fail to attend to, or to conceptualise, the shapes of its composing features (and the essential composition relations between the shape of the face and the shapes of the composing features).

Earlier in this chapter I suggested that the model of our perceptual experience of *objects*, propounded in chapter 2, provides us with the resources to see how perceptual experience might afford us with knowledge of the essence of an object in general (where 'object in general' is a sortal, denoting a very general natural kind). In chapter 2 we considered the tiger, padding through the Veldt, and tracking its prey. And yet, it seems that this tiger might have perceptual experience of objects relevantly like our own, and still might fail to think of objects in general at all. Just because it is possible to acquire knowledge of the essence of some property or natural kind through our perceptual experience of that property or natural kind, does not mean that doing so is inevitable. Plausibly, the tiger consistently fails to do this, despite the fact that her perceptual experience provides for it; plausibly, the tiger just does not have the right cognitive capacities to take advantage of everything that perceptual experience offers her.

Moreover, crucial to Allen's position on Revelation is an understanding of it on which it implies that perceptual experience affords *exhaustive* knowledge of essence. But nothing that I've said commits me to such a view. It seems quite plausible that perceptual experience provides us with some knowledge of the essences of the shapes and colours that we perceive, but that perceptual experience is neutral on, for example, the question whether there are any essential connections between the

colours and the (classes of) surface reflectance properties that they supervene on.

And it seems that Allen (2016) does in fact agree that perceptual experience affords us with knowledge of the essences of colours. Allen believes that the structural properties of the colours are essential to them, and he uses this claim to argue in favour of his preferred view of the metaphysics of colour – on which colours are distinct from the surface properties/reflectance types that they supervene on. Presumably this essentialist knowledge of the structural properties of colours is acquired through perceptual experience of colours.

After all, an alternative view that Allen (2016) considers and rejects, which seeks to deflate our essentialist intuitions regarding the structural properties of colours, claims that it is colour *experiences*, and not the colours themselves, that *essentially* instantiate the relevant structural properties. It seems, then, that the assumption in play here is that the origin of these (in their view, misleading) essentialist intuitions is our (in this case, colour) *experiences*.

Indeed, in footnote 7 of chapter 6, talking of a proposal from Davies (2014), Allen says: ‘it would still need to be shown that this approach can be adapted to allow for *experience* to provide knowledge of essential structural properties of the colours’ (italics added). Elsewhere, Allen says ‘careful reflection on colour experience might be necessary for knowledge of the essential nature of the colours’ (2016: 145). Clearly, then, Allen does believe that perceptual experience affords us with knowledge of the essences of colours.

If Allen is correct that this essentialist knowledge really does concern structural features of *mind-independent* properties (because, 1) they are structural features of the colours, and, 2), the colours are mind-independent) then he must believe that perceptual experience is capable of affording us with at least some knowledge of the essence of some mind-independent properties, consistent with his denial of Revelation. This is all that my argument requires. And endorsing this leaves us free to concur with Allen that ‘Acquaintance does not make knowledge of the essential nature of the colours improbably easy’ (2016: 154).

A second argument that Allen makes use of to argue that colours are *distinct* from the

surface properties/reflectance types that they supervene on has the same form as an argument that Kripke offered in favour of the claim that pain is *distinct* from the physical/neurological states that *it* supervenes on. He argues that we can perfectly well imagine worlds in which there are colours without the relevant surface properties/reflectance types, or in which the colours supervene on different surface properties/reflectance types (this second is a Naïve Realist gloss on the famous spectrum inversion thought experiment (Shoemaker: 1990, Block 1990)). He also argues that we can perfectly well imagine a world containing the relevant surface properties/reflectance types but in which there were no colours. Allen calls this his 'Modal Argument'.

The Modal Argument, Allen contends, supports the claim that the colours are distinct from the relevant surface properties/reflectance types. This is a claim about the *essence* of colour. But the above discussion of Chalmers, Yablo and Williamson, suggests that our capacity to imagine or conceive of these scenarios is only a guide to the essence of colour in so far as such imagining/conceiving is constrained by an accurate prior conception of the essence of colour. And this, I suggest, is provided by perceptual experience. So I propose that we view Allen's Modal Argument as making use of our faculty for counterfactual thought simply to make salient our pre-existing conception of the essence of colour, which is provided by perceptual experience. And this is perfectly consistent with Allen's emphatic rejection of Revelation.

6.4. Summary

In this chapter we considered what the proponent of the seeing-that only view might say about our knowledge of the (metaphysical) necessity of spatial composition relations. This led us to discuss our knowledge of (metaphysical) necessity more generally. I argued that conceivability/imagination-based accounts of such knowledge need supplementing with an explanation of where *essentialist* knowledge comes from. I suggested that, at least in some cases (including our knowledge of the essences of spatial properties), the source of our essentialist knowledge lies in perceptual experience, and, more specifically, in *acquaintance* with the relevant properties.

We also considered whether this explanation, in terms of acquaintance, might be co-opted by an Intentionalist like Tye, who does endorse the idea of an acquaintance relation. It was found that Tye's acquaintance relation is not a content fixing mechanism, and so is not meant to be a source of propositional knowledge. As such, on Tye's Intentionalist account, acquaintance cannot be the source of our essentialist knowledge. This means that if perception is to afford us with knowledge of essences, on Tye's account, then it must do so via its I-content. As we saw in chapter 1, there are two existing classes of account that seek to explain the fixation of I-content. These are Causal Accounts and Benefit-Based Accounts. In the next two chapters we shall consider each of these in turn, and I shall argue that both classes of account are incapable of explaining perceptual experience of the (metaphysical) necessity of spatial composition.

Chapter 7 – Causal Accounts of Content Fixation.

In chapter 1 I described and motivated the following thesis (Thesis 2):

Content-Securing Mechanism: Any notion of ‘content’ requires that there be some mechanism, consistent with that account, which explains how mind-independent properties and relations enter into the content of perceptual experience.

I there explained that the Naïve Realist has it that, in virtue of the acquaintance relation, the subject’s environment (more specifically, the portion of the environment which falls in her visual field; that portion which she is *acquainted* with) is what determines the content of her perceptual experience. This means that, for the Naïve Realist, the (relevant portion of the) subject’s environment is the content-securing mechanism. Because the Intentionalist allows that the content of perception can vary independently of the subject’s environment – i.e. it is possible to have a perceptual experience with the content ‘yellow banana directly ahead’ whether or not there is a yellow banana directly ahead -, she must provide an alternative content-securing mechanism. I also listed the 2 general forms that an Intentionalist mechanism of the required type might take. In this chapter we will consider whether causal relations might be invoked in order to explain how properties and relations get into the content of perceptual experience. I previously described causal accounts as follows:

Causal Accounts: A causal relation between types of perceptual states and mind-independent features of the environment gets those features into the content of the perceptual states.

In the next section I will examine some extant causal accounts of content fixation. I take it that the argument that I will offer, against such accounts, will apply to all causal accounts, indifferent of the details. But it will be worth getting clear on how the details of such an account might play out. To this end, in section 1 I will describe two

influential Causal Accounts of content fixation in some detail. In section 2, I lay out the Composition Objection. This objection makes use of Claim 2, established in chapter 4 (not to be confused with *thesis 2*, established in chapter 1):

- 2) Phenomenal-composition is a metaphysically necessary relation. It arises from the intrinsic nature of the composed and the composing phenomenal-shapes. The fact that phenomenal-composition arises from the intrinsic nature of the composed and the composing phenomenal-shapes explains why it is so easy to miss (and why it so far *has* been missed). Though it is a distinct element of phenomenal character, it is not something *extra* in the sense that it could be imaginatively subtracted from the composed and the composing phenomenal-shapes.

I will argue that Causal Accounts of content fixation are unable to accommodate perceptual experience of spatial composition relations in a way that is consistent with Claim 2, above. Then, in section 3, I consider three responses to the Composition Objection. The first two of these responses will prove unsuccessful, but the third will serve the Causal Theorists purposes. However, discussion of the Composition Objection will pave our way to another objection, which I take to be decisive. In section 4 I propound the, decisive, Modality Objection.

7.1. Causal Accounts of Content Fixation

Extant Causal accounts of content fixation have focused on *conceptual* content. Since the content of perception is generally thought to be (at least partly) non-conceptual (see chapter 1), this means that some of the issues that concerned the authors of these accounts may not be relevant when it comes the perceptual contents that we are interested in. Nonetheless, it will be worth getting a sense of the difficulties that have arisen for such accounts as they have in fact been formulated, and of the moves that have been made in an attempt to deal with these difficulties. As will become apparent, some of the more vexing difficulties for accounts of the fixation of *conceptual* content simply don't arise for an account of non-conceptual content. This might be expected to give the causal theorist some hope. However, as I shall argue in

subsequent sections, even causal accounts of the fixation of *non-conceptual* content meet an insoluble problem when faced with the composition/modal objections.

Rupert, talking of conceptual content, says:

'Relying on causal connections to fix term extension invites a difficult theoretical challenge, however. Of all of the causal relations into which a given language of thought term *t* enters, we must identify, in a principled way, those causal relations that fix the extension of *t*'. (1999: 322)

An obvious class of causal relations that need to be excluded are those involved in non-veridical experiences. If these were permitted to influence the content then non-veridical experiences would be impossible. Although Rupert is concerned specifically with the content of natural kind *concepts*, the point applies also to causal accounts of (probably non-conceptual) spatial perceptual contents. Representations possessing either type of content can be caused by a multitude of things. The various causal views of content differ in the way in which they propose to meet this challenge.

In fact, both Fodor and Rupert regard their proposals not only as accounts of the fixation of *conceptual* contents, but, specifically, conceptual contents apt to partake in something like a Fodorian language of thought. They focus on concepts of kinds of object/substance, because these present some problems that don't arise for representations of properties. In theory, then, the application of these accounts to non-conceptual contents pertaining to spatial properties ought to be more straightforward.

Fodor's (1987, 1990) account faces a number of challenges, but it is the most well known of the Causal Accounts of content fixation. He offers what he calls 'the asymmetric dependence' account of content fixation. According to this account it is a nomic relation that gets a property or kind into the content of some mental state. So, there is a law relating the mental representation HORSE to horses. However, not every horse causes tokens of HORSE, and not every token of HORSE is caused by a horse, so the nomic relations in question must be soft laws, where a 'soft law' is a law that has exceptions. It is to be expected, anyway, that when psychology (or any

'higher science') is the subject, any laws will be soft. The trouble is, soft laws are very easy to come by. If there is a soft law linking horses to HORSE, there will also be a soft law linking cow-on-a-dark-night to HORSE. And this is where Fodor's asymmetric dependence criterion comes in.

The above is Fodor's (1990) example. As Rupert (1999) points out, it is questionable whether cow-on-a-dark-night is an eligible candidate to appear in a law (soft or otherwise), for it is not a natural kind. In which case, the alternative candidate would simply be cows (at large), on account of the fact that cows sometimes cause tokens of HORSE on dark nights (and in other sub-optimal conditions). Of course, the temptation to say that there is a law connecting cows (at large) to HORSE may be reduced, relative to cow-on-a-dark-night, on account of the fact that the law is so unreliable. Given that HORSE is presumably a *conceptual* representation, the law connecting utterances of the natural language word 'horse' with the mental representation HORSE, is almost certainly stronger than the law connecting cows with HORSE. So this might be a better counterexample. One reason Fodor might be reluctant to concede, with Rupert, that only natural kinds may feature in laws, is that then his account would be powerless to explain the mental representations of artifacts.

Fodor's proposal is that wherever there are two candidates for the content of some representation, (unless the content is truly disjunctive, which, on all accounts, is the exception) the law linking one of the candidates with the mental representation will asymmetrically depend upon the law linking the other candidate with the mental representation. In which case, it is the law depended upon, and not the dependent law, that fixes the content. For example, the law linking cows-on-a-dark-night to HORSE asymmetrically depends upon the law linking horses to HORSE, so it is the latter law that fixes the content of HORSE. What this means is that, in the closest possible worlds in which the law connecting horses with HORSE is broken, the law connecting cows-on-a-dark-night with HORSE is also broken, but in the closest possible worlds in which the law connecting cows-on-a-dark-night to HORSE is broken, the law connecting horses with HORSE is not broken.

One might suppose that the intuition that Fodor is drawing on here is that in the closest possible worlds in which the law linking cows-on-a-dark-night to HORSE is broken, this is because our powers of discrimination are better than they in fact are. If this were why then one would expect the law linking horses with HORSE to remain unbroken. However, footnote 18 of Fodor's 1990 suggests otherwise. Talking of the application of his asymmetric dependence criterion to those representations that cause frogs to snap their tongues at flies, Fodor says:

'It's crucial that this claim be read *synchronically* since, presumably, frogs wouldn't develop a disposition to snap at black dots in worlds where the black dots have *never* been flies. The semantically relevant sort of asymmetric dependence is a relation among an organism's *current* dispositions. Take real-world frogs and put them in possible worlds where the black dots are bee-bees and they'll snap away, happy as the day is long. But real-world frogs in possible worlds where the flies aren't black dots are ipso facto snapless'.

If we're interested in the organism's *current* dispositions then the explanation for the failure of the law linking cows (on a dark night) to HORSE won't be that in that world we have superior powers of discrimination. What reason, then, do we have to assume that in such a case the law linking horses to HORSE is not also broken? Presumably the idea is that in that world cows look (even) less like horses than they do in the actual world, and this is why they are not mistaken for horses (even on dark nights). However, if this were right, wouldn't it also be the case that in the closest possible worlds in which the law linking horses to HORSE fails, this is because horses look different there? And why should this effect the law linking cows (on a dark night) to HORSE? After all, we're interested in the organism's *current* dispositions. We're instructed to take 'real-world' subjects and *put* them in possible worlds. As long as the cows in *that* world still look just as similar to the horses in *this* world, subjects will respond to cows (on dark nights) in that world just as they would respond to them in this world. But then the laws are *symmetrically independent*.

Maybe the idea is that, with time, the biologists among these people will establish that the strange horses there are, in fact, horses. Our *current* dispositions include the

acquisition of new dispositions – for example, the disposition to recognise strange horses as horses. So, the thought goes, following some scientific discoveries, the law linking horses to HORSE would be re-established (thanks to the endeavours of the scientific community). And recall that *all* horses there are strange. So the cows there look nothing like any of the horses there. It's plausible that once subjects have learnt to recognise the strange horses and have stopped expecting to see 'normal' horses, subjects there would lose the disposition to see cows (in any situation) as horses. This would mean that the law linking cows (on a dark night) to HORSES eventually fails, so there is an asymmetric dependence between the laws. In which case, we just have to let the situation play out a bit before we get the result that Fodor is after⁶⁹.

Now what of utterances of the natural language word 'horse'? The mental representation HORSE is about the biological species: horse. The mental representation is not about utterances of the word 'horse', though tokens of HORSE are frequently caused by such utterances (when the listener is an English speaker). Presumably Fodor's analysis would be that the law linking utterances of the typographically individuated word 'horse' with HORSE asymmetrically depends upon the law linking horses to HORSE, so it is the latter law that fixes the content of HORSE. But does it?

Paralleling what we did in the horse-cow case, we need to pick a possible world in which the relevant cause of HORSE – utterances of 'horse' - is altered, such that it is no longer a cause of HORSE (just as we altered cows such that they were no longer a cause of HORSE). Since utterances of any natural language words have the power to cause tokens of conceptual representations only in virtue of linguistic conventions, this time we shall have to put just a *single* 'real-world' subject in some possible world

⁶⁹ Would these people that we've *put* in this world come to think that the 'normal' looking horses had been transformed, or merely replaced? Williamson believes that we can, with some reliability, make use of counterfactual thought in order to derive philosophical knowledge (i.e. knowledge of identity conditions). But he acknowledges that it becomes more difficult to assess counterfactuals as these involve worlds that are more distant from the actual world; as, indeed, is usually the case when these are put to philosophical purposes (see chapter 6, and compare the kinds of counterfactuals that Williamson believes the capacity evolved to deal with).

that exemplifies an alternative socio-linguistic context⁷⁰. Then the Fodorian intuition might be that in the closest possible worlds in which the law linking utterances of the typographically individuated word 'horse' with HORSE is broken, this is because the word 'horse' means something different there, and some other natural language word there means horse (and, there, it is utterances of this other natural language word that frequently causes tokens of HORSE). As Fodor would hope, there is no reason to think that this would have any effect on the law linking horses to HORSE.

An issue presents itself. The relevant law links a type of mental representation, HORSE, to a typographically individuated type of word, 'horse'. It appears that this law will fail simply in virtue of the dispositions of everyone else in that world. But it seemed to be important, for Fodor, that it was the dispositions of 'real world' subjects that were implicated in the failure of the laws. So let's assume that the other subjects in this world don't possess the typographically individuated mental representation HORSE (though they will possess some other concept with the same content). In which case, the holding or failing of the law linking the word 'horses' to HORSE, in that world, will be determined entirely by the dispositions of our transplanted subject.

This helps, but given that the asymmetric dependence is to be understood as 'synchronic', this doesn't immediately deliver the right result. We are to assume that the subject has her *current* dispositions in the counterfactual world, so whether or not the typographically individuated word 'horse' means horses in that world, our subject's response to utterances of 'horse' will be the same. Of course, over time this will likely change. So perhaps, as with before, we just have to let the situation play out.

The trouble is, this doesn't seem like an asymmetric dependence. For it is not clear why we should think that in the closest possible worlds in which the horses-HORSE law fails, the 'horse'-HORSE law would also fail – in which case, we have a

⁷⁰ I must admit, I feel some discomfort here. As intimated in a previous footnote, it's difficult to know what the rules are when it comes to constructing what are fanciful counterfactuals. But let's give Fodor the benefit of the doubt. There are bigger problems to come.

symmetrical independence. There is no reason to think that people in the relevant world will speak a different language (after all, plausibly we are now to revert back to populating the possible world entirely with 'real-world' subjects, who will, of course, speak their 'real-world' language). Even if horses miraculously came to look different one day, and people in this strange world went through a period of failing, and then, perhaps, struggling to identify them, it's far from obvious that the word 'horse' would at any point lose the power to cause tokens of HORSE. In fact, it seems that horses would be the talk of the town ('Where have the horses gone?' 'Well I never, are *those* horses?'), so the law linking 'horses' to HORSE would be alive and well, despite the failure of the law linking horses to HORSE.

Moreover, it seems to me that, as Rupert (1999) suggests, the advocate of the Causal View *will* need to limit candidate contents to natural kinds. If Fodor does not limit candidate contents to natural kinds then it is possible to construct counter-examples. According to popular television program, QI, many breeds of sheep and goats look very similar (apparently the best way to tell them apart, in unclear cases, is by the direction in which their tails point). Take the class of all goats, and add a single sheep belonging to a breed that is difficult to distinguish from goats. This forms the class: goats-plus-this-particular-sheep. What is to stop this from counting as the content of GOAT? Since many goats look very much like sheep, and this particular sheep looks very much like a goat, this sheep may well even cause tokens of GOAT more reliably than the average goat does. In which case, the class consisting of all goats plus this particular sheep may well cause GOAT more reliably than will the class consisting only of all the goats. Significantly, it does not look like there is an asymmetric dependence between the law linking goats to GOAT, and the law linking goats-plus-this-particular-sheep to GOAT.

Call the special sheep (or, if QI is correct, the not so special sheep) that looks just like a goat, 'Shaun'. Even if there were a nearby world in which Shaun does not look the slightest like a goat, still the law linking goats-plus-Shaun to GOAT would seem to hold. It is a soft law (it must be), so the odd anomaly can be tolerated. And if there were a case for removing Shaun from the extension-determining set in order to improve the reliability of the law (in that world, at least), wouldn't there also be a

case to be made for removing from the set one of those goats that look an awful lot like sheep? Both moves improve the reliability of the law (in certain worlds), but only the former move brings the extension closer to the natural kind: goats.

Rupert (1999) sets forth his own proposal, which he calls the Best Test Theory (BTT). Unlike Fodor, Rupert explicitly restricts candidate contents to natural kinds. According to BTT the content of a mental representation is determined as follows:

BT1: If a subject S bears no extension-fixing intentions toward t , and t is an atomic natural kind term in S 's language of thought (i.e. not a compound of two or more other natural kind terms), then t has as its extension the members of natural kind K if and only if members of K are more efficient in their causing of t in S than are the members of any other natural kind.

Rupert believes that once the subject has at her disposal some basic range of concepts, the content of which is determined in accordance with BT1, she may then use these basic concepts in order to intentionally imbue further concepts with contents. The point of the first clause of BT1 is to make it clear that BT1 is designed only to explain the content of the basic concepts, which content must of course be settled before these concepts can then be used to imbue any further concepts with content. Since we are interested in the, probably non-conceptual, content of perceptual experience, BT1 ought to apply to the kinds of representations that we are interested in. This, then, is another respect in which the Intentionalist might hope that the fixation of *perceptual* content will be more straightforward than the fixation of conceptual content.

Key to understanding Rupert's proposal is his notion of being the most *efficient* cause of tokens of some concept. The most efficient cause of tokens of some concept is whichever natural kind has the highest *success rate* relative to that concept:

'Success rates are determined by the success rate function, $f(K, S, t, m)$. This function takes four arguments, one each from the following four categories: a natural kind (K), a subject (S), a natural kind term in that subject's language of

thought (t), and a time (m)... The output of the success rate function is determined by the ratio of the number of times members of K have caused a tokening of t in S to the number of times members of K have caused a tokening of any language of thought term in S . If, for example, members of K have caused the tokening of some term or other in S 's language of thought on 100 occasions, and 45 of the terms caused were tokens of t , then the success rate of K relative to S 's term t is 0.45. (323)

Some implications are worth spelling out. According to Rupert's view, HORSE could mean horse even though cows cause tokens of HORSE more frequently than horses do. This might be the case, say, if there are many more cows around than horses (and perhaps, for some reason, one is perpetually on the look out for horses, and thereby primed to mistake non-horses for horses). Still, according to Rupert's view, HORSE will mean horses because most occasions on which one sees a cow (and on which the cow causes, in one, a token of some concept) the cow will not cause a token of HORSE (but will instead cause a token of COW), so the success rate of cows relative to HORSE is relatively low. And presumably on almost all occasions on which one sees a horse (and on which the horse causes, in one, a token of some concept), rare though such occasions might be, the horse will cause a token of HORSE, so the success rate of horses relative to HORSE is very high indeed.

'Even if actual horses have only caused 5% of my 'horse' tokens, the success rate of horses relative to 'horse' may still be very high. For me, as well as for the typical subject, the success rate of horses relative to 'horse' is probably upwards of 99%... No other natural kind has a success rate relative to 'horse' in me which even approaches the success rate of the natural kind horse'. (324-325)

Rupert's view seems to commit him to denying that words are natural kinds. For otherwise it seems quite plausible that the word 'horse' might have a success rate relative to HORSE, in an English speaking subject, which is comparable to the success rate of horses in that subject. A difference between Fodor and Rupert's approaches that is worth noting is that while Fodor focuses on counterfactuals, Rupert focuses on the statistics of actual historical causal interactions (though in a footnote he does

acknowledge that 'circumstances may arise where it seems better to consider counterfactual causal interactions when assigning extensions to natural kind terms in the subject's language of thought' (330)).

All of the issues around natural kinds are less threatening when it comes to the, presumably non-conceptual, content of perceptual experience. For it is far from obvious that there are non-conceptual representations of artifacts as such. The restriction of eligible contents to natural kinds, then, is readily amenable to the Intentionalist who is interested in the non-conceptual content of perceptual experience. Moreover, there's no reason to think that tokens of non-conceptual visual representations are ever caused by utterances of words (e.g. 'horse'), or, (through association) by prior thoughts of anything else (e.g. the thought of saddles). So there is no worry about these kinds of things getting into the extension of the representations.

Of course, the general issue of fixing content remains. Tokens of perceptual representations have a chain of more and less proximate causes, and sometimes the distal cause of such tokens will be abnormal - as when we suffer an illusion and see what is red as green, or what is spherical as cuboid. So which set of these causes (many of which are natural kinds) is to determine the content of the representation? It is far from clear that there are non-conceptual perceptual representations of natural kinds like horses, or cows, so we'll continue to focus on perceptual experience of shape properties. In light of this, the example of seeing a sphere as a cuboid is the non-conceptual equivalent of the horse-cow case considered above. The Intentionalist's hope would be that one of these proposals, or some other causal account, can just show how to prevent the more proximate causes, as well as non-veridical perceptual experiences, from influencing the content of non-conceptual perceptual content (and thereby show how *non-veridical* perceptual experience is even possible).

Of the two accounts that we have considered, Rupert's seems the more apt for explaining the content fixation of spatial visual representations. This is because it's unclear what, given our *current* dispositions, would lead to a break in the law linking either spheres or cubes to SPHERE. It seems that we can't simply make the same

move that we made in the cow-horse case, because it seems that any change to the way that spheres or cubes look (where, again, this must be understood as a change in the properties, not in the dispositions of our perceptual systems) will constitute/require a change of the (type) identity of those properties⁷¹.

Over the course of this chapter it will transpire that if we adopt a causal account then, though we shall (just about) be able to explain perceptual experience of spatial composition, we certainly shall be unable to explain perceptual experience of the *necessity* of that composition relation in a way that is consistent with the phenomenology of visual experience. The reason for this is that, as we established in chapters 4-6, the necessity of the spatial composition relation manifests itself, perceptually, in the fact that phenomenal-composition arises from the intrinsic nature of the composed and the composing elements of phenomenal character. Assuming that this metaphysically necessary relation between the composed and the composing elements of phenomenal character is what is responsible for our perceptual experience of the necessity of spatial composition, causal accounts of content fixation are unable to explain perceptual experience of the necessity of spatial composition.

Before I explain the above argument, which I take to be conclusive, I will first consider how the Causal View handles perceptual experience of spatial composition.

⁷¹ Ryder (2004) presents a clever causal account of content, which, like Fodor's account, and for similar reasons, does not seem to have application to the contents that we are interested in. Ryder believes that a special class of neuronal cells - the pyramidal cells -, found in the cerebral cortex, *individually* represent types of substance. Ryder focuses on the representation of types of substance, e.g. banana, though he suggests that the account that he offers might have application also to types of event and even to particulars. He notes that it is rare that any individual property, or even set of properties, is a reliable guide to the type of some substance (or the type of some event, or the identity of some particular). The pyramidal cells attune themselves to the presence of types of substance by 'learning' to detect multiple sources of correlation - for example, a cell attunes to the presence of bananas by 'learning' to detect the correlation of yellow, brown or green, with the banana shape. Each cell 'learns' to compute some function, and the content of a given pyramidal cell firing is whatever caused the stabilization of the function which that cell computes (where this cause of the stabilization is some distal feature of the environment). This view does not have application to visual spatial property contents, for spatial properties do not seem to be sources of multiple correlations.

This will be the task of section 2. It will transpire that even this presents difficulties for the Causal Theorist. I will begin by presenting an argument for the claim that the Causal view is unable to account for perceptual experience of spatial composition. Then, in section 3, I will consider three responses that the causal theorist might propose. The first two of these responses will be unsuccessful, but the third will provide an out for the Causal Theorist. Nonetheless, this discussion will pave the way for the final, conclusive, objection to the Causal View, which I present in section 4.

7.2. The Composition Objection

If spatial composition enters into the content of perceptual experience then, given thesis 3 (below), there ought to be some element of phenomenal character that corresponds to the composition relation.

- 3. Phenomenal Spatial Contents:** A spatial property or relation enters into phenomenal content if and only if that property or relation enters into the content of a perceptual state, *and* there is an element of phenomenal character that uniquely corresponds to the perceptual content.

I have named this element of phenomenal character ‘phenomenal-composition’. Given thesis 3, on a causal account this element of phenomenal character has to be *selectively* causally responsive/sensitive to instantiations of the composition relation. Otherwise the dependence relation will not be unique in the way specified in thesis 3. If the element of phenomenal character is not *selectively* causally responsive/sensitive to instantiations of the composition relation – if some other property would also reliably cause the element of phenomenal character -, then the experience has a content that is indeterminate between the various properties that the phenomenal element is causally responsive/sensitive to.

This point is analogous to that which presented such problems for Causal Accounts of *conceptual* content, which we considered above. In that context, competing variants of the Causal View differed in the way that they attempted to filter out those causes that played a role in fixing the content of a representation from those that did not. As

I have already said, it seems that many of the issues that led to difficulties in that context simply will not arise in relation to non-conceptual content. But the above point also differs from the one that presented problems for Causal Views of *conceptual* content in an important respect. For now we are interested not in the causal relation between that which fixes content and the *vehicle of representation*, but, rather, in the causal relation between that which fixes content and the corresponding element of phenomenal character. Depending upon which dependence type the causal theorist espouses the element of phenomenal character may or may not be the *vehicle of representation* – only on dependence type b will it be so.

I lay out the premises of the argument below. The general argument is that, in virtue of the necessity of phenomenal-composition, the uniqueness of dependence relations, which is demanded by thesis 3, breaks down. Since the composing phenomenal properties necessitate the composed phenomenal properties, anything in the environment that can be said to cause the instantiation of phenomenal-composition will have to do so by causally influencing the composing phenomenal properties. Likewise, anything that causes the instantiation of a composing phenomenal property will at the same time cause instantiations of both the composed phenomenal property and also the phenomenal-composition relation between composed and composing phenomenal properties.

At this point how contents and phenomenal properties are paired becomes a matter of arbitrary decision. The pairing cannot actually be arbitrary, for this would be to give up on the idea that there must be a mechanism that determines phenomenal contents. So really we shall have to say that the composed phenomenal property, the composing phenomenal properties, and the phenomenal-composition relation, all correspond to a single phenomenal content, if they correspond to any at all. That single phenomenal content will be indeterminate between the three corresponding mind-independent properties – the composed shape, the composing shapes, and the composition relation between them. But we have already seen, in chapter 4, that there is much to be said in favour of the view that these are distinct phenomenal contents (though, in section 3 of this chapter, we shall encounter this view - that there is really just a single phenomenal content here - in a new guise).

Before I provide the formal argument for the claim that causal views of content cannot accommodate perceptual experience of spatial composition, I need to lay some groundwork. Specifically, I need to say a bit about the individuation of phenomenal-types, and corresponding phenomenal contents. There are more and less determinate ways of typing phenomenal-shape-properties. For example, one can count all phenomenal-straight-lines as a single phenomenal-type, and all instances of this phenomenal-type do correspond to a common phenomenal content.

There are also available more determinate typings of phenomenal properties, and these correspond to more determinate phenomenal contents. On a fully determinate typing of phenomenal-straight-lines they will be typed depending upon the length, orientation, and the location in the visual field, of the straight-lines that we suppose to be their usual causes (and, therefore, to be the objects of perception). So in order for some straight line to cause an instance of such a phenomenal-straight-line it will have to have a particular length, a particular orientation, and a particular location in the visual field. Nonetheless, this is a set of properties that, at different times, different straight lines can possess. There will, of course, be a huge number of maximally determinate phenomenal-straight-line-types. The argument that follows concerns maximally determinately typed phenomenal properties and phenomenal contents.

1. Perceptual content is fixed by causal relations between vehicles of representation and the properties that feature in those contents.

I.e., the causal view of perceptual content fixation. Suppose this is true.

2. Phenomenal Spatial Contents: A spatial property or relation enters into phenomenal content if and only if that property or relation enters into the content of a perceptual state, *and* there is an element of phenomenal character that uniquely corresponds to the perceptual content.

I argued for 2 in chapter 1 (it was there labelled 'thesis 3'). 'Correspondence' was fleshed out in the introduction as any of the following 4 possible dependence relations:

a. Each visual, spatial, phenomenal-type *uniquely* depends on one spatial perceptual experience content-type.

Or

b. Each spatial perceptual experience content-type *uniquely* depends on one visual, spatial, phenomenal-type.

Or

c. Each visual, spatial, phenomenal-type is identical with a spatial perceptual experience content-type.

Or

d. Each visual, spatial, phenomenal-type *uniquely* depends on one spatial perceptual representation vehicle-type, and each spatial perceptual experience content-type also *uniquely* depends on one spatial perceptual representation vehicle-type (from the very same set of spatial perceptual representation vehicle-types).

3. Spatial composition relations feature in the phenomenal content of perceptual experience.

This was one of the main conclusions that I argued for in chapters 4 and 5.

4. There is an element of phenomenal character corresponding to spatial composition relations: phenomenal-composition.

Follows from 2 and 3

5. Instances of phenomenal-composition must normally be (either directly or indirectly) caused by instances of spatial-composition.

Follows from 1 and 4

The causal relation will be direct if one espouses dependence type b, for on this view content is determined by phenomenal character, so the content-fixing causal relation must be between the property and the element of phenomenal character. The causal relation will be indirect on any of the remaining dependence types, for the content-fixing causal relation will be between the property and some representational vehicle, and the element of phenomenal character will either supervene on the vehicle (dependence type d), supervene on the content (dependence type a), or be identical with the content (dependence type c).

6. The phenomenal-composition relations that hold between different types of phenomenal-shape-properties belong *necessarily* to those types of phenomenal-shape-properties. The composed phenomenal-shape-property (the phenomenal-square, for example) has its source immediately in the intrinsic nature of the composing phenomenal-shape-properties (the maximally determinately typed phenomenal-straight-lines, for example).

I argued for 6 in chapter 4.

7. The causation of phenomenal-composition by spatial composition is screened off by the causation of composing phenomenal properties together with the necessitation relation between composing and composed phenomenal properties.

Follows from 6.

At this point the uniqueness of dependence relations (demanded by 2) breaks down.

8. Premise 5 is false: contradiction!

To avoid the contradiction we should reject premise 1 as all the other premises have been defended in previous chapters, or follow from those premises.

7.3. Responses

7.3.1 An Additional Element of Phenomenal Character

Perhaps the element of phenomenal character which we have been focusing on, and which we have established to be a *necessary* relation, is not really the element of phenomenal character that (uniquely) corresponds to spatial composition. In which case, what we have been talking about, and what we have been calling ‘phenomenal-composition’, is not really phenomenal composition at all. There may be some other element of phenomenal character actually deserving the title ‘phenomenal composition’; some feature which is not a relation that arises from the intrinsic nature of composed and composing phenomenal-shapes, and which is not vulnerable to the above screening off objection (premise 7). If this is correct then premise 6 of the above argument is false.

If we did find additional relations among the phenomenal-shape-properties, which did not arise from their intrinsic natures, then the Intentionalist could avoid the objection by saying that what we have been calling phenomenal-composition does not really correspond to any distinctive content at all (precisely because it is not ‘selectively causally responsive’ to any distinctive content). In other words, they could say that what we have been so far calling ‘phenomenal-composition-relations’ are epistemically redundant, because they are *merely* the product of the intrinsic nature of our phenomenal character. In a poetic turn of phrase, one could then call them ‘mere seams in the fabric of experience’.

Alternatively, the Causal Theorist could say that what we have been so far calling ‘phenomenal-composition’ is really mind-independent spatial composition. Our attention slips straight past the relevant elements of phenomenal character and onto the mind-independent relations that are represented. This second version of the response comes close to rendering phenomenal character something non-introspectable, in the spirit of Speaks (2015). But in the introduction, following Shoemaker (2000), we said that that if there is to be any point at all in talking about

phenomenal character then it must be something introspectable. So the idea must be that our attention often, though *not* inevitably, slips straight past the relevant elements of phenomenal character and onto the mind-independent relations that are represented.

In either case, the Causal Theorist must endorse a distinction between what we have been so far calling 'phenomenal-composition', and some additional element of phenomenal character, which we have so far missed, and which actually deserves the title 'phenomenal-composition'. The trouble with this response is that when we introspect the relations among phenomenal-shape-types, all we find is what we have all along been calling phenomenal-composition. Visual experience of spatial properties and their composition relations exhibits, then, a 'local' transparency. With regard to spatial properties and their composition relations, there seems only to be one set of properties and relations available to introspection (and these relations are those that we have been calling phenomenal-composition relations).

The present response supposes that there is some *additional* element of phenomenal character that does not have the modal properties of that which it serves to represent, and which we must surely think of *only* as an element of phenomenal character. It must not have the modal properties of that which it represents if it is to avoid the screening off objection, yet if it did not have those modal properties then it would surely (conspicuously) violate the 'local' transparency of our experience spatial properties and their composition relations. For then there would be *two* sets of properties available to introspection – those that represent and those that are represented, or those that represent and those that are epistemically redundant -, and these sets would be easily distinguished from one another by their *differing* modal properties. But this is not what we find.

The claim that there is not an additional element of phenomenal character that lacks the modal properties of that which it serves to represent is a 'local transparency claim', because it identifies a specific element of phenomenal character that is implicated by some view – the Causal View -, but which does not seem to be available to introspection. The implication is that what we have been calling 'phenomenal-

composition' must indeed be the element of phenomenal character responsible for getting spatial composition into the phenomenal content of perceptual experience, because introspection reveals no alternative.

7.3.2 The Analyticity of Spatial Relations

At this point it can be tempting for an advocate of the Causal View to reassert that the representation of the square *surely* does just follow analytically from the perceptual representation of the four located straight lines. My argument depends upon the claim that the content pertaining to the square does *not* follow analytically from the content pertaining to the four located straight lines. Otherwise there would be no need for separate elements of phenomenal character corresponding to the square and the four located straight lines, and so no need for a separate element of phenomenal character (phenomenal-composition) relating the square *to* the straight lines. This response, then, challenges premise 3, for it denies that composition relations feature in experience as genuinely *distinct* contents.

But we have already argued, in chapter 4, in favour of these distinctions. So what new grounds are there for doubt? On an absolutist conception of space, space is a system of relations that is capable of existing independently of any occupants. If someone has an absolutist conception of space then relations between (composed and composing) spatial properties will be grounded in the relations between their locations. Since I do not want to judge between relativist conceptions of space and absolutist conceptions of space, I will now consider an objection that arises if we assume an absolutist conception of space.

It has been suggested to me that my argument in this chapter relies on an implausible atomistic view of space. It is only if the square, the located straight lines, and the composition relation, are genuinely distinct (phenomenal) contents, requiring distinct elements of phenomenal character, that the screening off objection has application. But, the objection goes, since the nature of a given location is exhausted by that location's relations to other locations, it is incoherent to suppose that one could visually represent locations in the visual field without (at least implicitly) also

visually representing any relations between such locations. So, in virtue of these relations, which are necessarily (though perhaps implicitly) represented, the representation of the square does just follow analytically from the perceptual representation of the four located straight lines.

Let's suppose that 'a', 'b', 'c' and 'd' are the locations of four straight lines, which compose a square (so 'a' is a location co-extensive with some straight line at a time). If it is incoherent to suppose that location 'a' can be visually represented without visually representing its relations to, say, locations 'b', 'c', and 'd', then locations 'b', 'c', and 'd' must be at least implicit in the representation of 'a'. In which case, 'a', 'b', 'c', and 'd' are not logically distinct contents, and we ought not to expect distinct corresponding elements of phenomenal character. What's more, the shape described by 'a', 'b', 'c' and 'd' (and composed by the occupants of these locations) won't be a logically distinct content either, since this shape is the product of the relations between the four locations, and these relations are already implicit in the representation of 'a'. Clearly if this were correct then it would preclude the need for further representations (and corresponding elements of phenomenal character) pertaining to composition relations, and the screening off objection would fail.

Now, it is possible to deny that, strictly speaking, locations enter into the phenomenal content of perception at all. It is a coherent position to maintain that we experience spatial locations by experiencing the spatial properties of, and relations among, the objects that occupy those locations. But the present objection is premised on the assumption that locations do enter the phenomenal content of perception, and that the representations of these locations entail the various composition relations among the perceived shapes. So in responding I will assume that locations do feature in the phenomenal content of visual experience.

If my argument required that any region of space could be represented entirely independently of its relations to *any* other region of space then the above objection would present a real problem for me, but my argument does not require that. My argument requires only that each location in the visual field can be visually represented independently of every other location *in the visual field*. It is open to me

then to say that each visual-location-content entails a relation to the subject, or better, a relation to the subject's location/point of view (since this is *not*, generally, a *location in the visual field*). But why ought the objector to agree with this characterisation of visual location phenomenal contents? There are good reasons for thinking that, if locations do enter into the phenomenal content of perceptual experience, these locations must be *subject-relative* locations⁷². And since the location's relation to the subject (or the subject's point of view) is logically distinct from the location's relations to the rest of the perceived locations, then 'a', 'b', 'c', 'd', and the relations between these, must after all be distinct contents.

If someone kidnapped and blindfolded you, and took you to an empty white room, and then removed the blindfold, you would not know, simply on the basis of sight, where you were. Of course, you would be in a position to linguistically refer to your objective location – you can say, 'I am here' – but this is something that you could have done even if you had remained blindfolded. Blindfold on or off, you would not know *where* 'here' is. Suppose that you managed to get hold of one of your kidnapper's phones and you called for help. When asked your location you reply 'I am here'. This would not help because, still, neither you nor the person that you are speaking to would know to *where* 'here' referred. However, if we instead imagine that the room is furnished, you can know, purely on the basis of sight, at what subject-relative spatial location each of the room's contents is situated.

On the other hand, if spatial relations to the subject (or the subject's point of view) were not at least implicit in visual location content then even the entire set of relations among all locations in the visual field would not be sufficient to uniquely identify any set of objective locations. In such a case, all that would be visually represented is a system of relations. Moreover, assuming that the visual field describes a symmetrical space then if visual location contents did not entail any

⁷² Moreover, see chapter 7, of Price (2001), for an Intentionalist argument for the claim that the capacity to perceptually track a particular object requires the perceptual representation only of *subject-relative* locations. Price is a Consumer Semanticist, but, as Millikan (2007) points out, it seems that causal theories will anyway have trouble accounting for the representation of particulars (Ryder's (2004) causal account may be an exception – see footnote 70).

relations to the subject this would make it difficult to see how one could visually represent something as occupying one portion of the visual field rather than another.

The objection was that it is not possible to identify a location in the visual field without bringing in that location's relations to other locations in the visual field, so the location and its relation to other locations in the visual field cannot really be logically distinct contents. But we've found that the representation of relations between locations in the visual field will not serve to uniquely identify a set of locations, while the representation of relations with the point of view of the subject will serve to uniquely identify a set of locations. This point suffices to refute the objection. For the only relation that is entailed by the content of a subject-relative location representation is the relation between the subject and the location in question. It's then a further question what are the relations (if any) between that location and all of the other locations that the subject perceives. Once we admit that the location's relation to the subject is logically distinct from the location's relations to the rest of the perceived locations then the rest of my argument goes through. For if the relations between locations in the visual field each constitute distinct (phenomenal) contents then there will need to be distinct corresponding elements of phenomenal character.

7.3.3 Causal Over-Determination

This final response, which I take to be successful, challenges premise 7. More specifically, the response denies that premise 7 follows from premise 6. The response focuses on the notion of being 'selectively causally responsive'. It argues that the screening off objection trades on an unreasonably stringent understanding of what is required for a phenomenal property to be selectively causally responsive to some mind-independent property. The response comes at a cost, for it is not successful if the Intentionalist advocates dependence type c (on which phenomenal properties are identical to contents – see chapter 1), but it is successful for the remaining three dependence types.

In order for this response to work we must understand all occurrences of the word 'depends', in the above argument for the composition objection, in a particular way. Specifically, the word 'depends' appears several times in premise 1 - the theory neutral notion of 'phenomenal spatial contents' (which I first presented in chapter 1). In order to permit the advocate of the seeing-that only view adequate leeway to formulate the response we must condone a particularly loose reading of the word 'depends'.

On a strict sense of 'depends' it would be supposed that if c depends on b then there could be no c without b. For present purposes the seeing-thater requires a notion of 'dependence' on which that is not implied. If c depends on b then we can say that b *determines* c, or we can say that c occurs *in virtue of* b, but this explicitly leaves open the possibility that c might have occurred in virtue of something entirely different: a. This means that c can vary independently of b (and, by the same token, independently of a); b is sufficient for c, but not necessary for c. From now on, unless I state otherwise, when I talk of 'dependence' this is what I mean. This reading was already implicit in premise 1 (the theory neutral notion of phenomenal contents), since the alternative reading would have rendered redundant my use of the word 'uniquely'.

There are two steps to the present response. The first is to establish that phenomenal-composition can be causally responsive to spatial composition, despite the screening off objection. The second step is to argue that phenomenal-composition can be *selectively* causally responsive to spatial composition, despite the screening off objection.

The causal theorist can achieve the first step by maintaining that phenomenal-composition is causally over-determined. However, the Causal theorist needs to explain how spatial composition exerts any causal influence on phenomenal-composition at all. Given that phenomenal-composition arises from the intrinsic nature of composing phenomenal-shapes, the only way in which instances of spatial composition could cause instances of phenomenal-composition is if the instances of spatial composition causally altered the intrinsic nature of the instances of composing phenomenal-shape-properties. And spatial composition must be capable of causing

phenomenal-composition relations involving a disparate range of phenomenal-shapes (phenomenal-straight-lines, phenomenal-curves... etc.), so there must be some intrinsic feature common to all composing phenomenal-shapes, which spatial composition causes, and through which spatial composition affects phenomenal-composition. This common feature, the Causal Theorist might say, is simply that of being a phenomenally-composing phenomenal-shape (where this feature is a determinable property).

The second step is achieved by weakening the notion of what it is to be '*selectively* causally responsive'. The argument is that for phenomenal-composition to be '*selectively* causally responsive' to spatial composition does not require that phenomenal-composition *only* be causally responsive to one thing – i.e. to spatial composition. If it required that then the demand would be too strong, for everything is causally responsive to a multitude of things. Rather, the response goes, what is required is that phenomenal-composition be causally responsive to spatial composition – and this possibility has been secured by embracing causal overdetermination -, *and* that there be some non-arbitrary way of privileging the right cause (i.e. spatial composition) – and this second requirement is secured by the fact that spatial composition is more reliably (causally) connected to phenomenal-composition than any specific composing shape is.

The *relevant* dependence relation is with the composition relation, rather than with the composing shapes, on the basis that phenomenal-composition is *more* reliably caused by spatial composition than by any particular type of composing properties – for example, on those occasions when the composed property is a circle then phenomenal-composition will not be caused by any straight lines, but in such cases phenomenal-composition *will* still be caused by spatial composition.

This last response is not successful if the Causal Theorist espouses dependence type c (on which phenomenal properties are identical to contents – see chapter 1). The Causal View is a view about content fixation, so it is silent on the question of what makes something a representation in the first place. Generally speaking, Causal Theorists will say that it is either functional properties of a state (Fodor: 1987) or

teleological properties of a state (Neander: 2013, and, more tentatively, Fodor: 1990) that render that state a representation⁷³ (and then causal relations just fix the content of the representation). It might even be the case that whatever the phenomenal-square supervenes on is used (and evolved for the purpose of being used), by the visual system, as a *sub-personal* representation of a square. But if those representations are not also used at the *personal* level (or did not evolve for the purpose of being used at the personal level) - that is, if they are not used by the *subject* (or did not evolve for the purpose of being used by the subject) -, then the relevant content (i.e. square) does not feature in the subject's perceptual *experience*.

This implies that it is possible for a subject to visually represent the 4 straight lines without representing the square (or visa versa), and it is possible for the subject to visually represent both the 4 straight lines and the square without representing the composition relation between them. In order for all of these properties/relations to be visually represented, at the personal level, there must be states that causally covary with each of the properties, *and* which have the right functional or teleological properties. But we know that the 4 phenomenal-straight-lines necessitate the phenomenal-square (and the phenomenal-composition relation), so this implies that there can be phenomenal-squares (and phenomenal-composition) in the absence of any visual contents pertaining to squares (or to spatial-composition). This contradicts the assumption that these elements of phenomenal character are identical to the relevant contents, so dependence type c is false.

This might not seem like such a great hindrance to Intentionalism since dependence

⁷³ Fodor (1990) suggests that it is *all and only* representations that realise causal laws with the asymmetric dependence feature which he believes to be responsible for content fixation. If this is correct then his asymmetric dependence proposal might afford an account not only of content fixation but also of what makes something a representation in the first place. In which case, Fodor's (1990) might be capable of reconciling dependence type c with the above objection after all. However, he also there says that there may be a role for teleological properties in generating the *normative* dimension of representation. This makes Fodor's (1990) view sound more like Neander's (2013) view. Anyway, we have already seen that it is far from clear that Fodor's account has application to the contents that we are interested in. Moreover, an objection that has application to *all* causal accounts, regardless of dependence type, will be presented in the following section.

type c constitutes only half of the options available to the Representationalist – Representationalists may endorse either of dependence types c or a (see chapter 1) -, so the objection does not even serve to preclude any major branch of Intentionalism (i.e. Representationalism). Dependence type a does not identify elements of phenomenal character with contents – it instead has phenomenal properties *supervene* on contents -, and the view thereby creates some distance between phenomenal properties and contents, which is required in order to formally block the above objection.

However, as soon as one considers what the relationship between phenomenal properties and contents might actually be, on such accounts, dependence type a starts to look like less of a safe refuge for Representationalists. For the chances are that a Representationalist may well opt for dependence type a over dependence type c simply because, for example, they think that some element of phenomenal character, p, is identical *not* with some content, c, but with the mental state's *property of having that content*, c⁷⁴. In which case, the phenomenal character – the property of having content x - logically entails the relevant content. So although the phenomenal properties are not identical with their corresponding contents, there is not actually *enough* distance between the contents and the phenomenal properties to avoid the objection. For such an account would be unable to accommodate the possibility of the instantiation of phenomenal properties from the relevant class, without the tokening of representations with the corresponding contents. So the objection does in fact threaten Representationalism in general, and it should worry any Intentionalist who favours an externalism about phenomenal character.

The response to the composition objection, described in this section, does effectively deal with the screening off objection, provided that the Causal Theorist endorses one of the remaining dependence types – types b or d. And there at least remains the logical space for a (Representationalist) view that espouses dependence type a and that avoids the objection. But the Representationalist faces the challenge of

⁷⁴ See Fish (2010) for an explanation of Representationalism (or, as he calls it, 'strong content-first Intentionalism') on which this is the nature of the relationship between phenomenal properties and contents.

describing the relationship between phenomenal properties and contents in such a way that the phenomenal properties may be instantiated in the absence of the corresponding contents. We could call this the 'Freedom Constraint' – assuming a causal account of content fixation, on any viable Representationalist account the phenomenal properties must be permitted sufficient freedom to make the above response to the composition objection feasible. Most importantly, by going through the above three responses to the composition objection and by getting clearer about exactly which phenomenal properties must be responsible for our perceptual *experience* of spatial composition, we are now in a position to turn to what I take to be the fatal objection to (all versions of) the Causal View.

7.4. The Modality Objection

In chapter 4 we established that we not only perceptually experience spatial composition, but we seem to be perceptually sensitive to the profound intimacy of that relation. In chapter 4 I suggested that this sensitivity to the profound intimacy of spatial composition manifests itself in a corresponding phenomenal-intimacy (of, for example, phenomenal-squares and phenomenal-straight-lines), and I suggested that this phenomenal intimacy might explain why phenomenal-composition has until now gone unnoticed. Moreover, in chapter 4 we took the intimacy of phenomenal-composition to be a necessitation relation – phenomenal-squares necessitate phenomenal-straight-lines. We suggested that this, and our sensitivity to it, would explain our differing intuitions about the types of impossibility involved in two tasks: visualising a stadium full of faces replete with details, all at once (which is merely *empirically* impossible), and visualising a square without straight lines (which is *metaphysically* impossible).

Then, in chapter 6, I noted that the phenomenology, uncovered in chapter 4, is consistent with perceptual experience of the *necessity* of mind-independent spatial composition relations. I also offered an epistemological argument in favour of the claim that in virtue of the necessitation relation between phenomenal-straight-lines and phenomenal-squares we perceptually experience the necessitation relation

between composed and composing mind-independent shapes. So, we do perceptually experience spatial composition *as* a necessary relation.

It is the fact that perceptual experience of the *necessity* of spatial composition is achieved *in virtue* of the necessity of phenomenal-composition that creates insoluble difficulties for the Causal View. It requires that token composition relations among mind-independent shapes be the cause of the *necessity* of token composition relations among corresponding phenomenal properties. But, clearly, the very concept of a metaphysical necessity precludes a causal explanation. Something cannot be caused to be metaphysically necessary⁷⁵. At least, something cannot be caused to be metaphysically necessary by any spatio-temporal relation⁷⁶. To say that a relation can be *caused* to be metaphysically necessary is somewhat like saying that something non-living is non-living precisely because it was born. A birth, by its very nature, produces a life. Metaphysical necessity, by its very nature, is uncaused.

Even if, as was argued in section 3.3 of the present chapter, spatial composition can cause phenomenal-composition (and can count as *the* privileged cause of phenomenal-composition, and so qualify as *the* corresponding phenomenal content), still, spatial composition – more specifically, *the necessity* of spatial composition – cannot cause phenomenal-composition *to be* a necessary relation. So, for the Causal Theorist, it cannot be in virtue of the *necessity* of phenomenal composition – i.e. the phenomenal-intimacy that we uncovered in chapter 4 – that we perceptually experience the necessity of spatial composition.

The only move available to the Causal Theorist is to search for an alternative, additional element of phenomenal character, which serves to represent the necessity of spatial composition. This move is analogous to that made in section 3.1, in response to the screening off objection raised in section 2. And the move fails in the present case for very similar reasons, involving a ‘local transparency’.

⁷⁵ According to Williamson (2007) it is a principle of the modal logic S5 that ‘the necessary is necessarily necessary, and the possible necessarily possible’ (135).

⁷⁶ I will prescind from issues around the causal efficacy of God in relation to the metaphysically necessary.

As before, the Causal Theorist has two options. They may say that the nature of phenomenal-composition – the fact that it arises from the intrinsic nature of composed and composing phenomenal properties – is epistemically redundant. Or they may say that we are simply mistaking an observation regarding the necessity of spatial composition relations for an observation regarding the necessity of phenomenal-composition relations, because our attention tends to slip straight past the relevant elements of phenomenal character and onto the nature of the mind-independent relations that are represented.

But, in accordance with the transparency of experience, there just seems to be one set of properties and relations available to introspection. We may think of these properties and relations as mind-independent, or we may think of these properties and relations as elements of phenomenal character. The present response supposes that there is some *additional* element of phenomenal character that is not a modal property, but which serves to represent something (spatial composition) as having a specific modal property (as being necessary). This phenomenal property we must surely think of *only* as an element of phenomenal character. But when we introspect our experience, all we find is what we have been calling the ‘composing phenomenal properties’, the ‘composed phenomenal property’, and the ‘phenomenal-composition relation’ between them. Our awareness of the necessity of spatial composition has its source in our appreciation of the nature of *that* relation (the only relation available to introspection), not in some additional element of phenomenal character that represents the spatial composition relation as necessary (and which could have been absent, even while the nature of phenomenal-composition itself remained unchanged).

It has been suggested to me that it is not completely obvious that one metaphysically necessary relation couldn't cause another metaphysically necessary relation. Perhaps something can be caused to be metaphysically necessary so long as what's doing the causing is itself metaphysically necessary⁷⁷. Now, I'm not convinced by this response. But, anyhow, it is possible to construct an objection, of the same form

⁷⁷ Richardson: personal correspondence.

as that above but going in the opposite direction, which is not vulnerable to the present response. Can the *metaphysical* necessity of spatial composition (or of anything else) exert a causal influence on the visual system? It does not seem that a modal property of some relation is something that can be causally efficacious⁷⁸. This is precisely why the issue of modal knowledge – considered in chapter 6 – is so vexing. More specifically, this is precisely why, as Lowe (2012) and Roca-Royes (2011) explain, Williamson's (2007) attempt at an account of modal knowledge actually only accounts for knowledge of nomological (or empirical) necessity.

7.5. Conclusion

The main conclusion of this chapter is that Causal Accounts of content fixation are unable to accommodate perceptual experience of the *necessity* of spatial composition. However, there is an additional conclusion to be drawn, which does not rely on the claim that we perceptually experience the necessity of spatial composition.

In order to accommodate perceptual experience of spatial composition the Causal Theorist must loosen the connection between content and phenomenal character. The result of this is that existing Representationalist versions of the Causal Account are refuted. For it was found that, in order to accommodate perceptual experience of spatial composition, the Causal Theorist must reject dependence type c. It was noted that the Representationalist is still free to endorse dependence type a. However, we also noted that existing versions of dependence type a tend to construe phenomenal character as the property of having *such and such* content. So, although phenomenal character and content are not identified with one another on this view, existing construals of dependence type a still do not create *sufficient* space between contents and elements of phenomenal character in order that the Representationalist might accommodate perceptual experience of spatial composition.

⁷⁸ In which case, we have an argument against even *unconscious* perception of necessity on the Causal View.

Chapter 8 – Consumer Semantics and Benefit-Based Accounts of Content Fixation.

In section 1 I will offer a basic account of Consumer Semantics (CS). I will there offer an explanation, on the Consumer Semanticist's behalf, of how the Consumer Semanticist might offer an account of perceptual experience of spatial composition that is consistent with Claim 2, from chapter 4, but which, unlike Causal Views, does not require the postulation of causal over-determination (or the imposition of the Freedom Constraint). If this is possible then even Representationalist versions of CS will be able to deal with The Composition Objection.

The account of CS with which we start this chapter will be basic indeed. Once we begin to consider its application to the personal level representations employed by humans it will become apparent that the initial explanation of our perceptual experience of spatial composition will need to be complicated in various ways. Crucially, it is plausible that if perceptual experiences are representations then the primary use to which these representations are put in creatures like our selves is in the production of beliefs. This observation will cue a brief investigation of the content fixation of beliefs. In section 2 we will consider Price's (2000) CS account of the content fixation of (perceptual) beliefs, which are 'general-purpose representations'. We shall also consider perceptual experience in creatures, like us, who make use of such 'general-purpose representations'.

It will emerge from the discussion of section 2 that Price's (2000) CS account entails conceptualism. In section 3 I will draw out some of the counter-intuitive consequences of conceptualism. Then, in section 4, we shall apply both The Composition Objection and The Modal Objection to Price's account. It will be found that there is available to Price a response to the Composition Objection, similar in form to that utilised by the Causal Theorist in the previous chapter. But, ultimately, in making this response Price will have to postulate causal over-determination (and will have to accept the imposition of the Freedom Constraint). The implication is that

Representationalism is precluded by the response. Moreover, as in the previous chapter, the Modality Objection will prove decisive against *all* forms of the view under consideration.

8.1 – Consumer Semantics and Spatial Composition: The Basics.

Millikan (1984) proposes that the content of a representation is determined by the use to which some sub-personal representation consumer puts it. The idea is that the accuracy conditions of a representation are whatever must obtain if that representation is to enable the representation consumer to perform its proper function. The proper function of a representation consumer is then given a teleological reduction of the kind that can be given for other organs, such as hearts and lungs: *x* is the function of an organ *y* iff. *y* was selected/preserved because it did *x* in a 'critical mass of cases'.

A sub-personal representation consumer is a device that is designed to perform some invariant function, but which achieves this result by different means depending upon environmental factors. The representations are what inform the device of such environmental contingencies and enable the device to adapt itself to them, and to thereby perform its invariant function. For example, there are representations produced by the frog's visual system that are used by some mechanism, in the frog, to guide the frog's tongue snap behaviour. The (teleologically determined) invariant proper function of this tongue snap controller mechanism is to enable the frog to catch flies, we may suppose. It can do this by producing tongue snaps in a range of different directions (and at different times). So the representations, produced by the frog's visual system, are what inform the snap mechanism of relevant environmental contingencies and enable it to adapt itself to them. A given representation can be said to represent the presence of a fly in a certain location (at a certain time) because the representation modifies the activity of the snap mechanism such that it will only aid that mechanism in performing *its* proper function (catching flies) on the assumption that there is a fly at that location (at that time).

Neander (Stanford Encyclopaedia entry on Teleological Theories of Mental Content: 2012) offers a taxonomy on which Consumer Semantics is part of a broader class of accounts – benefit-based accounts - on which the content of some representation is whatever feature of the distal environment happens to be responsible for the fact that the representation *benefits* the organism (though I will call anyone who espouses a Benefit-Based account a ‘consumer semanticist’). On Consumer Semantics accounts this benefit-based explanation goes via the proper function of some sub-personal representation consumer. The representation adapts the sub-personal representation consumer to the presence of the relevant feature of the environment, such that the consumer is then able to successfully perform its proper function, thereby benefiting the organism. If the relevant distal feature of the environment did not exist then the representation would never successfully adapt the consumer to the environment, so the representation would never aid the consumer in performing its proper function (at least not in accordance with a normal explanation), and it would never benefit the organism (in accordance with a normal explanation). We shall see that sometimes it is necessary to formulate a benefit-based account without recourse to sub-personal representation consumers. This may be the case with ‘general purpose representations’. We shall return to these issues in section 2.

It seems like the representation consumers afford a way of dealing with the composition objection. The problem, as it arose for the Causal View, was that the uniqueness of dependence relations (demanded by thesis 3), between contents and phenomenal properties, broke down. Since the composing phenomenal properties necessitate the composed phenomenal properties, anything in the environment that can be said to cause the instantiation of phenomenal-composition will have to do so by causally influencing the composing phenomenal properties. Likewise, anything that causes the instantiation of a composing phenomenal property will at the same time cause instantiations of both the composed phenomenal property and the phenomenal-composition relation between these. At this point how contents and phenomenal properties are paired becomes a matter of arbitrary decision. But the pairing cannot actually be arbitrary, for this would be to give up on the idea that there must be a mechanism that determines phenomenal contents.

With representation consumers determining the content of a representation, it becomes possible to attribute unique dependence relations once again. Certainly, the four located straight lines might simultaneously cause the phenomenal-straight-lines, the phenomenal-square, and the phenomenal-composition relation, but since, on the present view, it is not causal relations (of this kind) that determine content, that is no problem. On the present view it is the use to which representation consumers put representations that determines their content.

If we assume phenomenal internalism, and if, for simplicity, we assume that elements of phenomenal character are (identical with) the vehicles of representation, we can say that it is the different uses to which representation consumers put the phenomenal-straight-lines, the phenomenal-square, and the phenomenal-composition relation, that endows these phenomenal properties with distinct contents. Each of these distinct contents depends *uniquely* upon only one of these three phenomenal properties, even though the phenomenal properties all necessitate one another, and even though they are all caused by the same set of properties in the environment. The causal theorist was also ultimately able to deal with the Composition Objection, but only by eschewing Representationalism. So, if Benefit-Based Accounts can deal with the Composition Objection, without precluding Representationalism, then they will provide the Representationalist with a response to that objection.

Perhaps the spherical red/green things in the subject's environment tend to be edible – they tend to be apples. In which case, there might be a representation consumer with the function of finding the subject food, which is using a representation to carry information about spheres, in conjunction with another to carry information about colour, in aid of performing its function. Perhaps there are also representation consumers that merely serve the purpose of facilitating object-directed behaviour. It will aid such representation consumers in performing their task if they have representations of the shapes that objects instantiate. For example, circular objects can be rolled. If the object is circular then a motor command that results in object rolling may be the best way for the representation consumer to fulfil its function. In

which case, recognition that an object is circular will be useful, for that representation consumer, and ultimately for the subject.

Likewise, a representation consumer might selectively use a representation for the purpose of carrying information about composition relations. Recall that we noted, in chapter 4, that we have the capacity to purposefully manipulate the shape of a ball of play dough with ease and apparent foresight. If we are asked to reproduce a target shape, we do not proceed in a trial and error fashion, until we happen upon the correct shape. Nor do we seem to have to recall shape transformation ‘rules’ that we have previously learnt. Further, we noted that perceptual experience of spatial composition would explain this capacity to simply look and copy. This suggests that the representation of some type of composition relation might have its content in virtue of being used by some representation consumer to enable it to perform its function of effecting shape manipulation.

However, the account of Consumer Semantics with which we started this chapter was basic indeed. Once we begin to consider its application to the personal level representations employed by humans it becomes apparent that the initial explanation of our perceptual experience of spatial composition needs to be complicated in various ways. Crucially, it is plausible that if perceptual experiences are representations then the primary use to which these representations are put in creatures like our selves is in the production of beliefs. This means that (an important class of) the sub-personal representation consumers that make use of perceptual representations would be mechanisms with the proper function of generating beliefs.

In section 2 we will consider the content fixation of beliefs – that is, general-purpose representations – and we shall consider perceptual experience in creatures, like us, who make use of general-purpose representations.

8.2 –Perceptual Content for Believers

There is a dearth of material, in the Benefit-Based tradition, on the content of human perceptual experience. However, Price, who has a highly developed Benefit-Based

account of human-type cognition, is a notable exception. Following Price, we begin with the following problem for benefit-based accounts. Talking of what she refers to as ‘teleological accounts’ of content, Price says:

‘Recently, a number of writers, including Kim Sterelny, John Campbell and Radu Bogdan have suggested that the theory has rather limited prospects. These writers are happy to concede that the theory works well when we are considering relatively simple intentional systems, such as the bee dance mechanism; but they suggest that it cannot account for more sophisticated intentional systems, such as our own doxastic system. The reason they give is that doxastic systems are *general-purpose* intentional systems; they suggest that the teleologist is unable, in principle, to give an account of general-purpose intentional content (Sterelny, 1990, p. 134; Bogdan, 1994, p. 181; Campbell, 1994, p. 212f; see also Fodor, 1990, p. 65f). In this paper, I would like to try to counter this suggestion by outlining a strategy which a teleologist might use to launch an account of general-purpose content’ (2000: 123).

Describing Millikan’s view of the distinction between special purpose and general purpose representations, Price says ‘One contrast she suggests is that, while less sophisticated states belong to systems that are dedicated to the satisfaction of a particular need or set of needs, our own beliefs and desires may be used to help to satisfy almost any need’ (2000: 124). And she quotes Millikan in this regard:

‘No strictures beyond relevance (some semblance of logic) determine which beliefs and desires may interact with which to form new beliefs and desires or help to produce actions; beliefs are not hooked to certain uses and unavailable for others. Contrast the toad’s belief that these are bugs, which is fixatedly hooked to its desire to eat bugs’. (Millikan, 1986, p. 72) (2000: 124)

Price considers the possibility of defining a special-purpose intentional system as ‘a system that is dedicated to satisfying some specifiable need, or perhaps some specifiable set of needs’ (2000: 124), but she identifies some problems with this definition. The definition relies on some antecedent individuation of needs.

‘An obvious suggestion is that we should equate an organism’s needs with its

basic biological needs: the need for food, the need to avoid predators, the need to reproduce, and so on. But this suggestion seems to push the problem one stage back, for it is far from clear how we should individuate biological needs. Indeed, it might be argued that both frogs and humans have only one basic biological need—the need to produce fertile and healthy offspring. If so, the idea that humans are able to represent their environment in a way that is independent of a particular biological need will collapse.’

For this reason, Price offers an alternative proposal:

‘Instead, I would like to focus on the patterns of behaviour that the organism is able to produce. Any organism that is capable of representing its environment will be able to engage in certain activities: snapping at flies, threatening rivals, signing cheques, and so on... The ability to engage in a certain pattern of behaviour could be said to generate certain subordinate needs, which I shall refer to as *interests*. For example, given that the frog normally satisfies its need for food by catching flies, we can say that the frog has an interest in locating flies... Patterns of behaviour, so construed, are individuated by their functions. An organism may produce a number of different patterns of behaviour, functionally defined, in order to satisfy what we would intuitively regard as a single basic need. For example, *moving towards prey*, *grasping prey* and *swallowing prey* might all be different patterns of behaviour that help to satisfy a predator’s need for food. For this reason, interests are sliced more thinly than basic biological needs. Nonetheless, patterns of behaviour cannot be cut infinitely finely: *moving halfway towards prey* may be something that the predator does, but it will not count as a pattern of behaviour because none of the bodily movements that the predator produces have the function to realize that behaviour. Hence we can insist that an organism has a determinate set of interests at any one time’ (2000: 124-125).

Finally, Price defines a general-purpose intentional system as:

‘A system that is not restricted with respect to the range of interests that it may come to serve’ (2000: 125).

But, in the interests of clarity, Price warns against conflating this distinction with Fodor's distinction between 'cognitive modules' and 'central systems':

'Fodor's distinction, as I understand it, focuses primarily on the processes of belief *fixation*: the idea is that cognitive modules are unable to draw on certain kinds of information in representing the world; in contrast, the distinction drawn here focuses on the way in which the representations produced by different intentional systems are used in guiding the organism's behaviour. If so, a special-purpose system need not be an informationally isolated module; nor can we rule out the possibility that an informationally isolated module might operate as a general-purpose intentional system' (2000: 125).

An implication of the proposal is that it will not be possible to attribute a representation with some content on the basis of the interest or set of interests that the representation serves (e.g. catching flies), for there is no specific (set of) interest(s).

Price proposes to explain the content fixation specifically of general-purpose representations that are learnt/acquired, and where these representations are used specifically in perceptual judgments. Price's account introduces a causal/informational component, so it can be seen as a hybrid of Benefit-Based and Causal Accounts. As we saw in the previous chapter, causal accounts face the challenge of finding a principled way of distinguishing, from the many causes implicated in the tokening of some representation, just those causes that serve to fix the content of the representation. Price's account draws on the use to which representations are put in order to answer this question, but without assuming that representations must serve some particular (set of) interest(s).

Before offering her account, Price first describes the mechanisms that she will appeal to. She starts by describing the learning mechanism responsible for the acquisition of those representations (or concepts) involved in some perceptual judgement:

'I take it that these processes of learning will have been governed by some learning mechanism ancillary to the doxastic system. Moreover, I take it that one of the functions of this learning system will be to ensure that when the

judgements produced by the system become involved in processes of inference they will do so in a way that is appropriate, given the information that they carry. In other words, this mechanism will function to produce a kind of *harmony* between the rules in accordance with which the doxastic system produces judgements and the rules in accordance with which it employs those judgements in reasoning about the environment.

I will label this kind of harmony *J-harmony*. The two sets of rules will J-harmonize with each other if the doxastic system is operating in a way that accords with the following principle: where J is a type of judgement produced by the system; L₁ is the set of rules governing the way in which judgements are used in inference, L₂ is the set of rules governing the way in which judgments are produced by the system, and C is some condition in the environment; then, if J, when used in accordance with L₁, tends to give rise to behaviour that would be appropriate if C obtained, then it will also be true that J, if produced in accordance with L₂, will carry the information that C obtains' (2000: 130-131).

Finally, Price arrives at the following view:

'We are now in a position to draw these three points together. Firstly, the doxastic system that produced J will normally assign to J a certain role in inference. At this point in the account, we do not need to specify what that role might be- we will simply say that J is normally assigned *some role or other*. Secondly, J or related judgments are produced in accordance with a set of rules that has ensured that judgments of this kind have *sometimes borne R* to trees or to tall things. Thirdly, the fact that J has been assigned this inferential role is *explained* by the fact that J or related judgments have *sometimes borne R* to trees or tall things; moreover, this explanation goes *via* the normal workings of a learning mechanism that has the function to ensure that the doxastic system operates in a J-harmonious way. What I would like to suggest is that it is (in part) in virtue of these facts that J normally carries the information that some tree is tall; and it is in virtue of this that J can be said to represent this information, and so to have the function to enter into inference in a way that is

appropriate to the presence of a tall tree' (2000: 131).

We are now in a position to consider content fixation for perceptual representations in subjects that, like us, are in possession of general-purpose representational systems.

'The functions performed by our perceptual subsystems are unlikely to bear much resemblance to the kinds of function performed by the sensorimotor systems possessed by simple organisms. This is because the former, unlike the latter, do not have direct control over the organism's behaviour. They control its behaviour only via the workings of the central, doxastic system' (Price, 2001: 197).

Perceptual experiences serve as *input* to a general-purpose representational system, and this does seem to preclude an account along the lines of those described in section 1.

'The function of the perceptual subsystems that serve a doxastic system is to provide information *for the doxastic system*. There is no way of specifying what kind of information these subsystems are supposed to provide until we know what kind of information the doxastic system normally uses' (Price: 2001, chapter 10, footnote 11).

'The representations produced by my face recognition system, for example, will play a highly restricted, stereotypical role in my psychology—for example, to prompt beliefs about the identity of the person that I am looking at' (Price, 2001: 197).

It seems, then, that the Consumer Semanticist ought to say that there are sub-personal mechanisms – call them perceptual judgement mechanisms (PJMs) – the proper function of which is to produce true beliefs, and that these mechanisms are aided in their performance of this function by perceptual experiences. The perceptual experience then has the function of co-operating with the PJM to produce a true belief. Perceptual experiences do this by mapping on to the environment according to some mapping function, as determined by the PJMs. The perceptual *system* has the

function of producing perceptual experiences that will co-operate with the PJMs to produce true beliefs, which means that it has the function of producing perceptual experiences that map onto the environment according to the relevant mapping function, as determined by the PJMs.

On this view, the representational significance of phenomenal-composition would then be secured by the fact that the PJMs respond to experiences involving phenomenal-composition by producing beliefs about spatial composition. Given that the PJMs respond in this way, perceptual experiences can only fulfill their function of co-operating with the PJMs to produce true beliefs if they map onto the environment in the relevant way. The content of the perceptual experience – the square is *composed* of 4 straight lines - is just the mapping function that must be preserved in order that the perceptual experience can fulfill this function. In this case the mapping function demands that there be phenomenal-composition if and only if there is spatial composition.

In the following section I will explore some implausible implications of this account. These are general implications, which have nothing in particular to do with perceptual experience of composition relations or modal properties. However, as well as highlighting some weaknesses of the view, an awareness of these commitments will be useful for framing our discussion of the Naïve Realist account of content fixation, in chapter 9. Then, in the final two sections of this chapter, I shall consider how Price's Benefit-Based Account fares against the Composition Objection and the Modal Objection. As with the Causal Account, it will be found that, with some fancy footwork, the Consumer Semanticist can (just about) respond to the Composition Objection. However, there will be no way for the Consumer Semanticist to deal with the Modal Objection.

8.3 – A Commitment to Conceptualism.

Crane (2009) distinguishes two ways of understanding 'conceptual content'. On one notion, content is conceptual only if it is composed of concepts. On another, more liberal, notion of 'conceptual content', a content is conceptual provided that one can only undergo an experience with that content if one possesses the concepts that one

would use in giving a verbal report of the content. It is in this second sense that Price's Benefit-Based account seems to be committed to conceptualism regarding perceptual experience. This is because, on this view, a perceptual experience is endowed with content in virtue of the use to which a perceptual judgement mechanism puts experiences of that kind. It is because there is a mechanism that is disposed to generate judgements regarding the colour red, in response to experiences of type x, that experiences of type x are possessed of content pertaining to the colour red⁷⁹. But the mechanism can only be disposed to generate judgements concerning the colour red if the subject is in possession of the concept 'red'.

An implication of conceptualism is that perceptual experience cannot outstrip our conceptual resources. This is a position that has been widely criticised. Colour experience seems to be a clear counter-example. Colour experience makes distinctions between shades that go well beyond our conceptual resources. We are able to see subtle differences between shades of colour, even when we lack names for the different shades, and when we are unable to re-identify instances of the same shade only a short time later.

The commitment to conceptualism has counter-intuitive implications regardless of which dependence type the Consumer Semanticist endorses. The implications, though, are different, depending upon whether the Consumer Semanticist is a phenomenal internalist (dependence types b and d) or a Representationalist (dependence types a and c). It seems that the problems generated by conceptualism are most acute if one is a Representationalist. Because Representationalists hold that phenomenal character is identical with, or supervenes on, representational content,

⁷⁹ This view comes perilously close to a view that I quickly rejected in chapter 1. I there said that it wouldn't do to say that an experience has a certain content because it is of a sort that disposes the subject to make judgments of a certain kind. The reason was that, assuming that the purpose of perceptual content is to justify the judgments that we make on the basis of those perceptual experiences, such an account would be circular. But circles become less threatening as they become bigger. This circle is a little bigger, and it may be just big enough that it is not vicious. For in this case it is not the judgments that endow the experiences with content, rather, it is the sub-personal mechanisms that generate such judgments that endow the experiences with content.

Representationalists must deny that the perceptual-type phenomenology is any more finely varied than our conceptual resources could permit. This seems to be flatly refuted by introspection, and also by our capacity for sorting and matching objects by their subtly different colours.

Since phenomenal internalists do not hold that phenomenal character supervenes on representational content then they can maintain that, just as it seems, our perceptual-type phenomenology *is* more finely varied than our conceptual resources would permit. But phenomenal internalists must deny that this extra variation in perceptual-type phenomenology has any epistemic import (so it is epistemically redundant). Accordingly, they must deny that, in virtue of the richer phenomenology, perceptual experience can offer any non-inferential warrant for judgements. So even though the subject undergoes a range of subtly different experiences, which systematically vary depending upon the shade that they are presented with, any judgements that subjects make in sorting and matching tasks will receive non-inferential warrant from the experience only if the subject has concepts for the specific shades in question.

Conceptualism also has implications regarding concept acquisition. Since there is no perceptual experience of the relevant property prior to acquisition of the corresponding concept, perceptual experience of some property cannot play a role in the acquisition of the corresponding concept.

The rationality of a subject's response to her perceptual experience depends upon how this response accords with the content of that perceptual experience. So, if a perceptual experience possesses a content pertaining to the property 'square', and a subject responds to this experience by acquiring the concept SQUARE, we can describe this as a rational response to the perceptual experience. But because, for the CS, perceptual experiences can have no such content *until* the subject has acquired the concept of a square, they are precluded from describing the acquisition of that concept as itself a rational response to perceptual experience. Moreover, the CS theorist can only say that the concept was 'learnt' (as opposed to *merely* acquired) if learning is not assumed to be a process subject to epistemic norms having their

source in *contentful perceptual experiences*. As we shall see, in the following chapter, the Naïve Realist *is* able to accommodate the role of perceptual experience in concept acquisition.

Regardless of which dependence type the Consumer Semanticist espouses, the commitment to conceptualism has counter-intuitive results. The issues presented, for the Representationalist, by conceptualism, are particularly acute, for in this context Representationalism has implications that are flatly refuted by introspection, and by our capacity for sorting and matching objects by their subtly different colours. We now return to the main argument.

8.4 – Price’s Benefit-Based Account, The Composition Objection, and The Modality Objection.

In this final section we shall apply both The Composition Objection and The Modal Objection to Price’s Benefit-Based Account. The way in which these objections apply to Price’s Benefit-Based Account is somewhat different from the way in which they apply to the Causal Account, though the results are much the same. In this case it is the conceptualism implied by the account that produces the problematic results.

The Composition Objection is that the Intentionalist is unable to account for perceptual experience of spatial composition relations in a way that is consistent with the phenomenology. It will be found that there is available to the Consumer Semanticist a response to the Composition Objection, similar in form to that utilised by the Causal Theorist in the previous chapter. However, as in the previous chapter, the response will be unavailable to the Representationalist. The cumulative effect of these counterpart arguments, in chapters 7 and 8, is that existing version of Representationalism are refuted, *without* reliance on the claim that we perceptually experience the necessity of spatial composition.

We shall also apply The Modal Objection to Price’s Benefit-Based account. The Modal Objection is that the Intentionalist is unable to account for perceptual experience of the *necessity* of spatial composition relations in a way that is consistent with the

phenomenology. And, as in the previous chapter, it is the Modal Objection that will prove decisive against *all* forms of the view under consideration (Representationalist or otherwise).

Perhaps a subject has the concepts of composed shapes, like squares, but does not have the concepts of straight lines or of spatial composition (see chapter 5). Assuming that she undergoes experience like our own (and not of some alien kind – see chapter 4), then if she is disposed to generate beliefs regarding squares on the basis of her perceptual experience, her experience must involve phenomenal-squares. In which case, given that phenomenal-squares necessitate phenomenal-straight-lines and phenomenal-composition (see chapter 4), the subject's experience must also involve phenomenal-straight-lines and phenomenal-composition. Yet, in such a case, these additional elements of phenomenal character would not constitute phenomenal *contents* – they would be epistemically redundant⁸⁰.

If the subject were to acquire the concepts of straight lines and of spatial composition, and if she were to begin making judgements, using these concepts, in response to her visual experiences, then the phenomenal-straight-lines and the phenomenal-composition (which both *already* characterise her visual experience) would at that point come have epistemic significance. That is, those elements of phenomenal character would then constitute phenomenal *contents*.

So, the subject can visually represent the 4 straight lines without visually representing the square (or *visa versa*), and it is possible for the subject to visually represent both the 4 straight lines and the square without visually representing the composition relation between them. But we know that the 4 phenomenal-straight-lines necessitate the phenomenal-square (and the phenomenal-composition relation), so this implies that there can be phenomenal-squares (and phenomenal-composition) in the absence of any personal level visual contents pertaining to squares (or to

⁸⁰ I offer this now only as an observation that will, momentarily, enable us to refute Representationalist versions of Price's account. However, in the final chapter the above observation, along with other observations made in the previous chapter, will be developed into another objection – The Harmony Objection –, which will apply to all non-Representationalist forms of Intentionalism.

spatial-composition). This is inconsistent with dependence type c, since it contradicts the assumption that these elements of phenomenal character are identical with the relevant contents, so dependence type c is false.

And, as we saw in the previous chapter, once dependence type c is refuted on grounds of these kinds, dependence type a starts to look untenable. The tenability of a dependence type a view will depend upon the specifics of the proposed relationship between the phenomenal properties and the corresponding contents. For example, the Representationalist may think that some element of phenomenal character, p, is identical *not* with some content, c, but with the mental state's *property of having that content*, c₈₁. In which case, the phenomenal character – the property of having content x – logically entails the relevant content. So although the phenomenal properties are not identical with their corresponding contents, there is not actually *enough* distance between the contents and the phenomenal properties to avoid the objection. The implication is that the objection threatens Representationalism in general, and it should worry any Intentionalist who favours an externalism about phenomenal character. As with Representationalist versions of the Causal view, Representationalist versions of Consumer Semantics will be viable only if they can meet the 'freedom constraint' (see chapter 7).

For the rest of this chapter I will take it that dependence type c has been refuted. Since there at least remains the logical space for a dependence type a account, in what follows I will not assume that dependence type a has been refuted. As we shall see, there is further trouble in store even for those who espouse dependence types a, b or d.

Recall that:

Phenomenal Spatial Contents: A spatial property or relation enters into phenomenal content if and only if that property or relation enters into the

81 See Fish (2010) for an explanation of Representationalism (or, as he calls it, 'strong content-first Intentionalism') on which this is the nature of the relationship between phenomenal properties and contents.

content of a perceptual state, *and* there is an element of phenomenal character that uniquely corresponds to the perceptual content.

If one is a phenomenal internalist, and one endorses dependence type b (on which content depends upon phenomenal properties – see chapter 1), then there appears to be no trouble in accommodating the above demand. Given the argument from chapter 7, we know that composing spatial properties will cause instances of phenomenal-composition, and so one might suggest that this causal role for the composing properties screens off the spatial composition relation from exerting any causal influence on phenomenal-composition. But since, on the present view, such causal relations do not determine content, this does not present a problem. The unique correspondence between contents and phenomenal properties is achieved provided only that some representation consumer makes *selective* use of the relevant phenomenal property.

With regard to Representationalists there is available a screening off objection of the kind considered in chapter 7. This is because the representationalist believes that phenomenal properties depend upon (dependence type a) their corresponding contents (for we ruled out the other version of representationalism, on which phenomenal properties are identical with contents – dependence type c - above). Since the visual representation of the four located straight lines suffices for phenomenal-straight-lines, a phenomenal-square, and the phenomenal-composition-relation (because of the necessitation relation between these) there appears to be no unique correspondence relation between these phenomenal properties and the contents that they are meant to correspond to.

The representationalist can respond to the objection in the same way that the Causal Theorist responded to this objection in chapter 7. Assuming that the subject possesses all of the relevant concepts, there will be a systematic over-determination of phenomenal properties by visual contents. If a subject undergoes an experience involving a phenomenal-square then she will also undergo an experience involving phenomenal-straight-lines and phenomenal-composition. The representation of the 4 straight lines will suffice for phenomenal-composition. But this does not preclude the

representation of the spatial composition relation from also being sufficient for phenomenal-composition. In which case, phenomenal-composition is causally over-determined. The Representationalist can accept this, but maintain that the *privileged* cause of phenomenal-composition is the representation of the spatial composition relation (not the representation of the 4 located straight lines), because this most reliably causes phenomenal-composition (see chapter 7 for a more detailed explanation of this response). So, phenomenal-composition does *uniquely* correspond to the content pertaining to spatial composition.

A similar objection can also be raised against phenomenal internalists who endorse dependence type d (on which both phenomenal properties and contents depend upon some class of representational vehicles – see chapter 1). This is because this theorist must maintain that the unique correspondence between contents and phenomenal properties is achieved because some representation consumer makes *selective* use of the vehicle that the phenomenal property *uniquely* supervenes on. However, given that an appropriate set of 4 phenomenal-straight-lines necessitate a phenomenal-square (and phenomenal-composition), whatever representational vehicle the phenomenal-straight-lines supervene on will also be sufficient for the phenomenal-square (and for the phenomenal-composition). So, in this case, it doesn't seem that there is a *unique* correspondence relation between representational vehicles and phenomenal properties.

The advocate of type d dependence can again respond that certain phenomenal properties are over-determined by representational vehicles. For example, two representational vehicles are sufficient for phenomenal-composition: the vehicle that we should like to associate with phenomenal-composition, but also the vehicle that we should like to associate with the phenomenal-straight-lines. But, despite this over-determination, we can still associate the representational vehicles with those phenomenal properties that we'd like to. For the *privileged* cause of phenomenal-composition is the representational vehicle that is used to represent the spatial composition relation (not the representational vehicle that is used to represent the 4 located straight lines), because this representational vehicle most reliably causes phenomenal-composition (see chapter 7 for a more detailed explanation of this

response). So phenomenal-composition does *uniquely* correspond to the content pertaining to spatial composition.

Now we turn to the conclusive objection. The Consumer Semanticist seems unable to explain how we perceptually experience spatial composition *as* a necessary relation. Key to this argument is the claim, already established, that we perceptually experience the necessity of spatial composition in virtue of the necessity of phenomenal-composition.

If the Consumer Semanticist is a Representationalist then they must maintain that a visual content pertaining to the necessity of spatial composition causes/subvenes (dependence type a) the necessity of phenomenal-composition (for dependence type c has already been refuted). But the content cannot cause/subvene the necessity of phenomenal-composition for the kinds of reasons that were offered in chapter 7. I argued that the metaphysical necessity of some token relation couldn't depend upon anything because such dependence contradicts the very notion of metaphysical necessity. This line of argument was put under some pressure, at the end of the chapter, by the suggestion that perhaps something can be caused to be metaphysically necessary so long as what's doing the causing is itself metaphysically necessary. Though I did not capitulate to the pressure, rather than resisting it I instead offered another (albeit formally similar) objection, which was invulnerable to the pressure. However, in the present case the response never gets any traction, for what's doing the causing is a *contingent* spatio-temporal event (i.e. the tokening of some content), and something *certainly* cannot be caused to be metaphysically necessary by a *contingent*, spatio-temporal event.

If the Consumer Semanticist is a phenomenal internalist then they must maintain that some representation consumer either uses the necessity of phenomenal-composition to represent the necessity of spatial composition (dependence type b) or uses some state that the necessity of phenomenal-composition supervenes on for that purpose (dependence type d). The latter option is ruled out for reasons already considered: the necessity of phenomenal-composition cannot depend upon any spatio-temporal event, so it cannot supervene on any mental state. The former option requires that

the necessity of phenomenal-composition be causally efficacious, for no sub-personal representation consumer could make use of the necessity of phenomenal-composition if this were not causally efficacious. But it does not seem that a modal property of some relation is something that can be causally efficacious. This is precisely why the issue of modal knowledge – considered in chapter 6 – is so vexing. More specifically, this is precisely why, as Lowe (2012) and Roca-Royes (2011) explain, Williamson’s (2007) attempt at an account of modal knowledge actually only accounts for knowledge of nomological (or empirical) necessity.

8.5 Conclusion

The main conclusion of this chapter is that Consumer Semantics/Benefit-Based Accounts of content fixation are unable to accommodate perceptual experience of the *necessity* of spatial composition. However, there is an additional conclusion to be drawn, which does not rely on the claim that we perceptually experience the necessity of spatial composition.

In order to accommodate perceptual experience of spatial composition the Consumer Semanticist must loosen the connection between content and phenomenal character. The result of this is that existing Representationalist versions of CS are refuted. For it was found that, in order to accommodate perceptual experience of spatial composition, the Consumer Semanticist must reject dependence type c. The Representationalist is still free to endorse dependence type a. However, as we noted in chapter 7, existing versions of dependence type a tend to construe phenomenal character as the property of having *such and such* content. So, although phenomenal character and content are not identified with one another on this view, existing construals of dependence type a still do not create *sufficient* space between contents and elements of phenomenal character in order that the Representationalist might accommodate perceptual experience of spatial composition.

The cumulative effect of the above argument, and of its analogue in the previous chapter, is that no existing version of Representationalism can accommodate

perceptual experience of spatial composition. And this holds even if one denies that we perceptually experience the necessity of spatial composition.

It is striking that, despite the difference between the ways in which these arguments apply to the two classes of accounts – Causal Accounts and Benefit-Based Accounts -, the results are extremely similar. Though it is not possible to assess whether or not any future account of content fixation might deal with these objections, arguably, the fact that all existing accounts produce such similar results gives us reason to suspect that the issue is with Intentionalism in general.

Chapter 9 - Naïve Realist Accounts of Content Fixation.

The Naïve Realist need appeal to no clever theories to explain our perceptual experience of spatial composition relations, or to explain perceptual experience of the necessity of those relations. The account is very simple, and for many the account will seem unsatisfying, for it is not a reductive explanation. If I were to rest content with demonstrating only that Naïve Realism can explain our perceptual experience of composition relations and their necessity then the task would hardly deserve a chapter of its own. But there is more to say.

In section 1 I will explain how the Naïve Realist accommodates perceptual experience of spatial composition, and of its necessity. In section 2 I explain how, from the foregoing discussion, we can extract an additional argument in favour of Naïve Realism. This additional argument does not depend upon the claim that we perceptually experience composition relations as necessary.

9.1. The Naïve Realist Account of Perceptual Experience of Spatial Composition, and of its Metaphysical Necessity.

The Naïve Realist believes that an element of phenomenal character associated with perceptual experience of some property is token identical with that property. So the Naïve Realist accommodates perceptual experience of spatial composition by saying that phenomenal-spatial-composition is identical with spatial composition simpliciter. Moreover, the necessity of phenomenal-spatial-composition is identical with the necessity of spatial composition simpliciter. For the Naïve Realist, it is a brute fact that when we are acquainted with shapes we are also acquainted with the composition relations among them. Likewise, for the Naïve Realist, it is a brute fact that when we are acquainted with composition relations we are acquainted with the necessity of those relations.

It is the desire for more illuminating explanations that drives the reductionist's project in all domains. In philosophy of perception the reductionist attempts to

understand the acquaintance relation in more basic terms – in terms of representation. In offering such an account it is incumbent upon the reductionist that they explain, in terms of representation, why we are acquainted with what we are acquainted with. And it is reasonable, honourable even, to pursue such explanations. The desire for explanations is part of what drives all philosophy. But when it is shown, in some domain, that all such reductive accounts fail, it is not reasonable to go on demanding such explanations.

Nonetheless, it might still be possible for the Naïve Realist to offer some kind of systematisation of which properties a subject will be acquainted with. When one perceptually experiences a square and a triangle, certain of their composition relations, in virtue of which the square and triangle are similar, necessarily make for a *phenomenal* similarity.

This makes some crude principle, like the following, tempting:

N) If we perceptually experience a property p, any property that is entailed by the intrinsic nature of p will be manifest phenomenologically.

But an implication of such a principle would be that we perceptually see a square as having the negative property of not being a triangle, for example. And this does not seem plausible. However, the focus on essences, in chapter 4, suggests a reformulation of the above principle that would rule out negative properties:

E) If we perceptually experience a property p, any essential property of p will be manifest phenomenologically.

As Lowe (2012) and Fine (1994) explain, there are many more metaphysically necessary truths than essential truths. The (merely) metaphysically necessary truths hold in virtue of the essential truths. Lowe offered the following example of the difference:

‘(Ei) An ellipse is the locus of a point moving continuously in a plane in such a fashion that the sum of the distances between it and two other fixed points

remains constant.

(E2) An ellipse is the closed curve of intersection between a cone and a plane cutting it at an oblique angle to its axis greater than that of the cone's side'.
(2012: 936)

According to Lowe, only (Ei) is an essential truth. Of (E2) Lowe said:

'This metaphysically necessary truth holds in virtue of the essences of an ellipse and a cone, which are two quite distinct essences. It is because of what an ellipse is, and what a cone is, that this relationship necessarily holds between ellipses and cones. But it is not part of anything's essence that it holds' (2012: 939)

For the same reason, the statement 'that square is not a triangle' is a metaphysically necessary truth, but not an essential truth. For it is because of what a square is, and what a triangle is, that this relationship necessarily holds between squares and triangles. It is not part of anything's essence that it holds.

This helps, but there remain some problem cases. For example, it might follow from the above principle that we perceptually experience the triangle as having angles that add up to 180 degrees, and this might be thought implausible. After all, we are not able to tell, simply by looking (i.e. without the requisite knowledge), that a triangle has angles that add up to 180 degrees. Moreover, the above principle might entail that we perceptually experience all manner of complicated mathematical/geometrical properties. When mathematical/geometrical properties are described using units like 'degrees', as in the property of having more than 170 degrees, the conventional status of these units gives us a principled reason for excluding the properties as candidate perceptual contents. This is significant, for it seems that any complicated mathematical/geometrical property will be expressed with the use of some unit.

This would not preclude the mind-independent properties themselves (which, being mind-independent, are also unit-independent) from manifesting themselves in perceptual experience. But it seems that the Naïve Realist ought to accept that

they are. What prompted us to deny this was the fact that we are unable to tell, simply by looking (i.e. without the requisite knowledge), that a triangle has angles that add up to 180 degrees. But perhaps this is simply due to the unit-free nature of our perceptual experience of this property. This is supported by the fact that, once equipped with a protractor, which provides the relevant knowledge concerning units, we *are* able to tell, by looking, that the sum of the angles of some perceived shape is 180 degrees (and one can do this even if one does not have the background knowledge that this is true of all triangles, or if one did not recognise that this particular shape was a triangle).

Admittedly, one cannot tell that a triangle has angles that add up to 180 degrees *simply* by looking (even when equipped with a protractor), for one must add the degrees of the three angles together. But one does not *infer*, on the basis of one's perceptual experience, that the triangle has angles that add up to 180 degrees. So, arguably, the requirement that we add the degrees of the three angles together reflects a limitation in our visual *access* to the relevant unit-free property, rather than a limitation in our perceptual content.

We saw in chapter 2 that the content of perceptual experience outstrips what we actually access, and can even outstrip what we are capable of accessing. Campbell described a hypothetical tiger that used colour to select an object even though it was unable to access the colours of those objects, and, as Campbell pointed out, this seems to be the actual situation for human infants. For adult human subjects, who are able to access colours, this is such an easy task that it can be surprising to hear that infants struggle with it. But even adult humans, equipped with a protractor, are unable to access the sum of the angles of a perceived shape with the ease with which they have learnt to access the colours of perceived objects. Access to *this* property remains forever (relatively) arduous. But this is a limitation of our visual *access* to the property, not a limitation in our perceptual content.

However, if we wish to maintain that there are some essential truths concerning the properties that we perceive, which perceptual experience simply cannot

provide us with knowledge of⁸², then E) is still too strong. When we are acquainted with a property we may not be acquainted with all of that property's (essential) properties. But this is ought not to surprise us. Nobody would expect it to be the case that whenever we are acquainted with an object we are acquainted with all of that object's (essential) properties. For example, Kripke (1980) famously suggested the essentiality of origins, for particulars. Suppose that Kripke is correct. No one would expect it to be the case that whenever we are acquainted with a particular we are also acquainted with that particular's origins. In which case, we may have to say that it is simply a brute fact that, when we perceive some property, certain of its essential properties manifest themselves in perceptual experience, but others do not. Once again, the Naïve Realist is not offering a reductive account, so they are under no obligation to offer an explanation of why, when we are acquainted with some property, we are acquainted with certain of that property's essential properties, but not with others. Perhaps the refusal to offer a reductive explanation puts the burden of proof on the Naïve Realist, but if so, this is a burden that has now been discharged⁸³.

In the next section I explain how, from the foregoing discussion, we can extract an additional argument in favour of Naïve Realism, over Intentionalism. This additional

⁸² Allen (2016) seems sympathetic to this view.

⁸³ Stoneham & Brewer (personal communication) have applied some pressure at this point in the dialectic. They argue that if one believes perceptual experiences to be a part of the natural order – i.e. to have causes and effects – then one is under an obligation to explain what causes one to be acquainted with exactly those properties that one is acquainted with. If, as I have argued, the metaphysical necessity of spatial composition is not causally efficacious then it looks like the Naïve Realist will be in no better position to explain this than are any of the alternative accounts. If this line of argument is successful, then the preservation of the claims that I have made in chapters 4-6 will come at the cost of having to maintain that, when we're acquainted with a property, we are acquainted with *all* of that property's *essential* properties. The discussion of this section suggests that this is not an absurd position to hold. When considering our perceptual experience of angles, we found that it is possible to accommodate perceptual experience of a surprisingly expansive range of essential properties, consistent with the manifest limitations on our capacity to make use of such perceptual contents. Moreover, the counter-example that I offered to this idea – which traded on Kripke's essentiality of origins, for particulars -, might also be consistent with this, if the claim is restricted only to property types, and not to particulars.

argument is important because it does not even depend upon the claim that we perceptually experience composition relations as necessary.

9.2. A Further Advantage of Naïve Realism: Harmony of Content and Phenomenal Character.

In the previous two chapters I argued that, in order to accommodate perceptual experience of spatial composition (though not with experience of its necessity), the Causal View and CS must loosen the dependence relation between phenomenal character and content such that experiences might sometimes involve phenomenal-composition even though spatial composition does not feature in the content of those experiences. In such cases there would be a loss of harmony of content and phenomenal character. Recall, in chapter 1 we established that there must be *some* dependence relation between content types and phenomenal types. We outlined 4 possible forms this dependence relation might take, for the Intentionalist. One of these forms was:

- c) Each visual, spatial, phenomenal-type is identical with a spatial perceptual experience content-type.

In order to deal with the composition objection advocates of both the Causal View and CS were forced to reject dependence type c (and probably dependence type a too, so (at least) existing forms of Representationalism were thereby refuted).

In chapter 7 we noted that:

‘The Causal View is a view about content fixation, so it is silent on the question of what makes something a representation in the first place. Generally speaking, Causal Theorists will say that it is either functional properties of a state (Fodor: 1987) or teleological properties of a state (Neander: 2013, and, more tentatively, Fodor: 1990) that render that state a representation (and then causal relations just fix the content of the representation). It might even be the case that whatever the phenomenal-square supervenes on is used (and

evolved for the purpose of being used), by the visual system, as a *sub-personal* representation. But if those representations are not also used at the *personal* level (or did not evolve for the purpose of being used at the personal level) - that is, if they are not (or did not evolve for the purpose of being) used by the *subject* -, then the relevant content does not feature in the subject's perceptual *experience*.

This implies that it is possible for a subject to visually represent the 4 straight lines without visually representing the square (or *visa versa*), and it is possible for the subject to visually represent both the 4 straight lines and the square without visually representing the composition relation between them. In order for all of these properties/relations to be visually represented, at the personal level, there must be states that causally co-vary with each of the properties, *and* which have the right functional or teleological properties. But we know that the 4 phenomenal-straight-lines necessitate the phenomenal-square (and the phenomenal-composition relation), so this implies that there can be phenomenal-squares (and phenomenal-composition) in the absence of any personal level visual contents pertaining to squares (or to spatial-composition). This contradicts the assumption that these elements of phenomenal character are identical to the relevant contents, so dependence type c is false'.

As a result of the above possibility the Causal Theorist must loosen the connection between phenomenal character and content (so they are forced to reject dependence type c, and probably dependence type a too). The relationship between phenomenal properties and contents must be such that phenomenal properties may be instantiated in the absence of the corresponding contents - we called this the 'Freedom Constraint'. Accordingly, the Causal View renders possible experiences with the right phenomenal character structure, and the right correlational properties, to be perceptual experiences of some property - spatial composition -, yet which, according to the theory, lacks some further property putatively necessary for being perceptual experiences of that property. But when we consider the particular case our intuition, surely, is that these further properties are not, after all, necessary for

the experience to count as a perceptual experience of a square (or of spatial composition)⁸⁴.

This loosening of the connection between phenomenal character and content makes possible a certain loss of harmony between content and phenomenal character. In the case that we considered, harmony was lost because the phenomenal character of spatial experience had a structure that mirrored that exemplified by the spatial properties in virtue of spatial composition, yet despite this, and despite the obtaining of the relevant correlations with relevant features of the environment (squares and spatial composition relations), the view denies that these experiences count as perceptual experiences of squares (or spatial composition relations).

In chapter 8 we found that, in order to deal with the composition objection, CS also had to loosen the connection between phenomenal character and content, and reject dependence type c (and probably dependence type a too). So CS is also required to meet the Freedom Constraint. But the reasons for this, in the case of CS, are slightly different. The implication of Price's CS is that, assuming that the concept of a square is not innate (and even if it is innate, surely there is still a *possible* subject who does not possess this concept), perceptual experience could exhibit phenomenal-squares (and phenomenal-composition) even when the subject does not (yet) possess the concept of a square (or of spatial composition). But in such a case phenomenal-squares will not (yet) represent squares. Phenomenal-squares only comes to represent squares once the subject acquires the concept of a square, and once some belief-generating

⁸⁴ The Phenomenal Intentionalist might have a way of avoiding this objection. The Phenomenal Intentionalist endorses dependence type b, and additionally believes that there is a narrow perceptual content that is determined exhaustively by phenomenal character (which is also assumed to be narrow). The Phenomenal Intentionalist might claim that not only the content of a representation, but also the attitude (i.e. that it is a *perceptual* representation, and, therefore, that it is a *representation*) is determined by phenomenal character. Phenomenal Intentionalists, Chalmers (2006) and Thompson (2010), admirably postulate wide perceptual content in addition to the narrow content, and also a causal mechanism to explain the fixation of wide perceptual content. It's difficult to imagine what equivalent explanation might be offered for *how* phenomenal features determine something a perceptual representation. And if such explanations cannot be offered then what is the point of reducing acquaintance to representation?

mechanism is thereby in a position to make the relevant use of the subject's perceptual states. So the CS theorist, also, is forced to make moves that make possible a certain loss of harmony between content and phenomenal character.

CS, like the Causal View, renders possible experiences with the right phenomenal character structure, and the right correlational (informational) properties, to be perceptual experiences of some property – spatial composition –, yet which, according to the theory, lacks some further property putatively necessary for being perceptual experiences of that property. But, once again, when we consider the particular case our intuition surely is that these further properties are not, after all, necessary for the experience to count as a perceptual experience of spatial composition.

The possibility of this loss in harmony of content and phenomenal character, implicated by both the Causal View and CS, is one that arises even if one denies that we perceptually experience spatial composition relations as necessary. Naïve Realism has no trouble with such cases, for, given that phenomenal character and NR-content are identical, harmony of content and phenomenal character is guaranteed on this approach. As such, Naïve Realism better accords with our intuitions about the content of these hypothetical experiences.

This objection has some features in common with traditional Swampman worries about causal and teleological views of content fixation. Swampman was imagined by Davidson (1987):

'Suppose lightning strikes a dead tree in a swamp; I am standing nearby. My body is reduced to its elements, while entirely by coincidence (and out of different molecules) the tree is turned into my physical replica. My replica, The Swampman, moves exactly as I did; according to its nature it departs the swamp, encounters and seems to recognize my friends, and appears to return their greetings in English. It moves into house my and seems to write articles on radical interpretation. No one can tell the difference.

But there *is* a difference. My replica can't recognize my friends; it can't

recognize anything, since it never cognized anything in the first place ...I don't see how my replica can be said to mean anything by the sounds it makes, nor to have any thoughts' (443-444).

Davidson's analysis of Swampman – as having no intentionality – holds for any version of the Causal View that appeals to *actual* causal relations (as opposed to counter-factual causal relations), and for any view of content fixation that appeals to teleological properties. The implication, in the present context, is that because Swampman has not, historically, been embedded within, and interacting with, a compositional, spatial world (i.e. has not been among the putative contents of her perceptual states), her perceptual states cannot have content pertaining to the spatial world. Many have thought this result counter-intuitive⁸⁵.

The Swampman case also interacts with one's view of phenomenal character. It has been suggested that it would be particularly counter-intuitive if Swampman underwent experiences with phenomenal character like our own (or any phenomenal character, for that matter), yet was to be denied intentionality. So some have thought that one way to ameliorate the force of Swampman, as an objection to Causal and Teleological views, is to espouse Representationalism (for then Swampman would also be denied phenomenal character, on account of being denied intentionality). Given that we ruled out Representationalist versions of both Causal and Benefit-Based Views in the previous two chapters (that is, *without* reliance on the claim that we perceptually experience the *necessity* of spatial composition), this response has now been definitely precluded. The Naïve Realist, on the other hand, is free to attribute to Swampman undiminished perceptual experience, if that is what the Naïve Realist's intuitions dictate. For the Naïve Realist does not propose any role for historical (causal/teleological) properties in securing perceptual content⁸⁶. So the arguments in the previous two chapters, against Representationalism, *do* enhance

⁸⁵ Though see Millikan (1996) for a defiant response. She even considers the possibility that we might be moved to deny Swampman human rights. So you could say that she 'bites the bullet'.

⁸⁶ Of course, problems may remain if the Naïve Realist holds that perception does involve representations, or that other areas of cognition involve representations, especially where these representations are thought to ordinarily be associated with phenomenal character.

this old objection to Teleological accounts, and to certain Causal accounts, of content fixation.

However, the primary purpose of bringing up Swampman now is to help to elucidate the Harmony Objection, which I have been elaborating in this section. The Harmony Objection differs from Swampman in three important respects. The first is that in the present example the subject *has*, historically, been embedded within, and interacting with, a compositional, spatial world (i.e. has been among the putative contents of her perceptual states). And in the present example, unlike Swampman, the subject has a long evolutionary history within that same environment.

The second difference is that our subject is not envisaged to be a physical, or even a functional, duplicate of some normal person (Davidson, for example). But the subject is envisaged to be *almost* physically and functionally identical to any one of us. The difference between the subject and any of us is only that the subject is not (currently) disposed to use phenomenal-squares, or whatever these supervene on, *as* representations. Nor does she have an evolutionary lineage that imputes to phenomenal-squares, or to whatever these supervene on, the *function* of being used, by the subject, *as* representations. This is because her ancestors never used phenomenal-squares, or whatever these supervene on, *as* representations.

The third difference is that in this case the subject is not supposed to be globally non-intentional. This means that we can raise questions about the epistemic value/significance of phenomenal-composition for a subject who is, uncontroversially, otherwise intentional. So we can perfectly well imagine what it might be like to be such a subject, occupying such a point of view, and we can ask whether or not an experience involving a phenomenal-square (and phenomenal-composition), in such a case, would suffice for some kind of a mental representation of a square (and of spatial composition). And, intuitively, it would do so.

Of course, the Intentionalist may maintain, if the phenomenal-square (or whatever this supervenes on), is not used by the subject as a *perceptual* representation (or didn't evolve for that purpose), then the phenomenology of the experience will not be

exactly the same as that associated with a *perceptual* representation of a square (just as they may well maintain that *visualising* a square might be phenomenally different from *perceiving* a square, because of functional or teleological differences between these representations). So the hypothetical experience would not be, phenomenally, exactly the same as a *perceptual* representation of a square.

But we are being asked to imagine an experience that involves a phenomenal-square and yet which implicates *no* personal level representation of a square whatsoever (perceptual or otherwise)! And the objection is that it is difficult to imagine how any experience involving a phenomenal-square might fail to possess any content (perceptual or otherwise) pertaining to a square. Because the Naïve Realist identifies personal-level perceptual content with phenomenal character, the Naïve Realist, unlike the Intentionalist, can maintain that a phenomenal-square is sufficient for personal-level content pertaining to a square.

And actually, given that the 4 phenomenal-straight-lines, out of which the phenomenal-square is composed, do implicate *perceptual* content, it *is* difficult to imagine how the phenomenal-square can implicate anything but a *perceptual* content pertaining to a square. And, again, because the Naïve Realist *identifies* personal-level perceptual content with phenomenal character, the Naïve Realist, unlike the Intentionalist, can maintain that a phenomenal-square is sufficient for personal-level *perceptual* content pertaining to a square (indeed, for the Naïve Realist, the phenomenal-square is sufficient even for the object of perception – the *mind-independent* square). Visualisation will then be given an analysis similar to the Naïve Realist analysis of hallucinations, which we considered briefly in the Introduction, in connection with Martin (2002): visualisation, like hallucination, is an experience that merely *seems* to possess those phenomenal features that would be sufficient for its object (in this case, a mind-independent square). So when one visualises a square, there really is no phenomenal-square at all.

Thesis Conclusion

We began with the observation that an intense and protracted discussion – 30 years of discussion – of the putative transparency of experience (TE) has proved it incapable of adjudicating between the competing accounts of perceptual experience. The feature of TE that this discussion has focused upon, and which participants in the discussion had hoped would be decisive in favouring their own preferred account of perceptual experience, was the manifest mind-independence of that which shows up when we introspect our experience. But the claim that what shows up to introspection is *manifestly* mind-independent was found to be contentious.

However, we noted another feature of TE that is less contentious. Namely, that when we introspect our experience, the properties that we find *appear* to belong to objects. In chapter 2 I explained why the Naïve Realist might have a problem in accounting for this phenomenon. Although the Naïve Realist builds the object into their characterisation of perceptual experience, right at the outset, this fact does not explain how it is that objects manifest themselves in perceptual experience *as such*. That is, it doesn't explain why it *seems* to the subject, when she introspects, as though she is experiencing an object. In chapters 2 and 3, drawing on Campbell's (2014) work, I offered an original Naïve Realist account of this phenomenon.

In chapter 4 I described an original phenomenological observation. This feature of phenomenology I called 'phenomenal-composition'. I argued that phenomenal-composition is the manifestation, in perceptual experience, of spatial composition. I also argued that our phenomenology is consistent with the claim that we perceptually experience spatial composition relations *as* metaphysically necessary. Then in chapters 5 and 6 I offered additional, epistemological, arguments in favour of the claims that we perceptually experience spatial composition relations, and that we perceptually experience them *as* necessary. I argued that all alternative accounts of our *knowledge* of these things were, at best, highly implausible.

In the final third of the thesis I used what I had established in chapters 4-6 to motivate Naïve Realism over Intentionalism. Chapters 7-9 each addressed different accounts of content fixation. In chapters 7 and 8 I argued that both of those classes of such accounts that are available to the Intentionalist – Causal Accounts and Benefit-Based Accounts – are unable to account for our perceptual experience of the necessity of spatial composition relations in a way that is consistent with the phenomenological observations made in chapter 4. In both of these chapters I made use of what I called a ‘local transparency argument’. The transparency claims that these arguments depended upon, being ‘local’, are invulnerable to objections of the kind advanced by Block (1996) and Kind (2003), which we considered in the introduction. Moreover, these transparency claims made no contentious appeal to the *manifest* mind-independence of that which presents itself to introspection.

In addition to the main objection to Intentionalism, which relied on the idea that we perceptually experience the necessity of spatial composition relations, there were some additional arguments that did not rely on this claim. We established in chapters 7 and 8 that the Causal View and CS both permit a certain loss in harmony of content and phenomenal character. Both views make possible experiences with the right phenomenal character structure, and the right correlational properties, to be perceptual experiences of some property – spatial composition -, yet which, according to the theory, lack some further property putatively necessary for being perceptual experiences of that property. The objection is that when we consider the particular cases our intuition is, surely, that these further properties are not, after all, necessary for the experience to count as a perceptual experience of spatial composition. In chapter 9 we found that the Naïve Realist is able to deny the possibility of any such loss in harmony of content and phenomenal character. As such, Naïve Realism better accords with our intuitions about the content of these hypothetical experiences.

Another implication of this loosening of the connection between content and phenomenal character, which both the Causal Theorist and the CS theorist were forced to make in order to accommodate perceptual experience of spatial composition, was to refute existing versions of Representationalism. It was found that, in order to loosen the connection between content and phenomenal character,

and to accommodate perceptual experience of spatial composition, the Intentionalist must reject dependence type c. It was noted that the Representationalist is still free to endorse dependence type a. However, we also noted that existing versions of dependence type a tend to construe phenomenal character as the property of having *such and such* content. So, although phenomenal character and content are not identified with one another on this view, existing construals of dependence type a still do not create *sufficient* space between contents and elements of phenomenal character in order that the Representationalist might accommodate perceptual experience of spatial composition. Again, *this* argument does not rely on the claim that we perceptually experience composition relations as necessary.

As mentioned at the end of chapter 8, it is striking that, despite the differences between the ways in which these arguments apply to the two classes of existing Intentionalist accounts – Causal Accounts and Benefit-Based Accounts -, the results are extremely similar. Though it is not possible to assess whether or not any future account of I-content fixation might deal with these objections, arguably, the fact that all existing accounts produce such similar results gives us reason to suspect that the issue is with Intentionalism in general.

References

- Allen, K., 2016. *A naive realist theory of colour*. Oxford University Press.
- Armstrong, D.M., 1989. *Universals. An Opinionated Introduction*. Boulder: Westview.
- Baylis, G.C. and Driver, J., 1993. Visual attention and objects: evidence for hierarchical coding of location. *Journal of Experimental Psychology: Human Perception and Performance*, 19(3), p.451.
- Bayne, T., 2009. Perception and the reach of phenomenal content. *The Philosophical Quarterly*, 59(236), pp.385-404.
- Block, N., 1990. Inverted earth. *Philosophical perspectives*, 4, pp.53-79.
- Block, N., 1996. Mental paint and mental latex. *Philosophical issues*, 7, pp.19-49.
- Campbell, J., 2002. *Reference and consciousness*. OUP Oxford.
- Campbell, J., 2005. Transparency versus revelation in color perception. *Philosophical Topics*, 33(1), pp.105-115.
- Campbell, J. and Cassam, Q., 2014. *Berkeley's Puzzle: What Does Experience Teach Us?*. OUP Oxford.
- Chalmers, D., 2002. Does conceivability entail possibility?. *Conceivability and possibility*, pp.145-200.
- Chalmers, D.J., 2005. *The Matrix as metaphysics* (pp. 132-176). na.

Chalmers, D., 2006. Perception and the Fall from Eden. *Perceptual experience*, pp.49-125.

Cohen, M.A., Cavanagh, P., Chun, M.M. and Nakayama, K., 2012. The attentional requirements of consciousness. *Trends in cognitive sciences*, 16(8), pp.411-417.

Crane, T., 2009. Is perception a propositional attitude?. *The Philosophical Quarterly*, 59(236), pp.452-469.

Davidson, D., 1987. Knowing one's own mind. *Proceedings and Addresses of the American Philosophical Association* 60 (3):441-458.

Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General*, 113, 501-517.

Egly, R., Driver, J. and Rafal, R.D., 1994. Shifting visual attention between objects and locations: evidence from normal and parietal lesion subjects. *Journal of Experimental Psychology: General*, 123(2), p.161.

Farkas, K., 2008. Phenomenal intentionality without compromise. *The Monist*, 91(2), pp.273-293.

Fine, K., 1994. Essence and modality: The second philosophical perspectives lecture. *Philosophical perspectives*, 8, pp.1-16.

Fish, W., 2010. *Philosophy of perception: A contemporary introduction*. Routledge.

Fodor, J.A., 1987. *Psychosemantics: The problem of meaning in the philosophy of mind* (Vol. 2). MIT press.

Fodor, J.A., 1990. *A theory of content and other essays*. The MIT press.

- Fodor, J.A., 1998. *Concepts: Where cognitive science went wrong*. Oxford University Press.
- Goldman, A.I., 1967. A causal theory of knowing. *The journal of Philosophy*, 64(12), pp.357-372.
- Green, E.J., 2017. A layered view of shape perception. *The British Journal for the Philosophy of Science*, 68(2), pp.355-387.
- Harman, G., 1990. The intrinsic quality of experience. *Philosophical perspectives*, 4, pp.31-52.
- Heil, J., 2003. *From an ontological point of view*. Oxford University Press on Demand.
- Hillstrom, A.P. and Yantis, S., 1994. Visual motion and attentional capture. *Perception & psychophysics*, 55(4), pp.399-411.
- Horgan, T. and Tienson, J., 2002. The Intentionality of Phenomenology and the Phenomenology of Intentionality. *Philosophy of mind: classical and contemporary readings*. Oxford University Press.
- Horgan, T., Tienson, J. and Graham, G., 2004. Phenomenal Intentionality and the Brain in a Vat. *The externalist challenge*, pp.297-317.
- Huang, L., 2010. What is the unit of visual attention? Object for selection, but Boolean map for access. *Journal of Experimental Psychology: General*, 139(1), p.162.
- Huang, L. and Pashler, H., 2007. A Boolean map theory of visual attention. *Psychological review*, 114(3), p.599.
- Johnston, M., 1992. How to Speak of the Colors. *Philosophical studies*, 68(3), pp.221-263.

Kahneman, D., Treisman, A. and Gibbs, B.J., 1992. The reviewing of object files: Object-specific integration of information. *Cognitive psychology*, 24(2), pp.175-219.

Kind, A., 2003. What's so transparent about transparency?. *Philosophical Studies*, 115(3), pp.225-244.

Kripke, S. A., 1980. *Naming and necessity*. Harvard University Press.

Lewis, D., 1980. Veridical hallucination and prosthetic vision. *Australasian journal of philosophy*, 58(3), pp.239-249.

Lowe, E.J., 2012. What is the source of our knowledge of modal truths?. *Mind*, 121(484), pp.919-950.

Martin, C.B., 2007. *The mind in nature*. Oxford University Press.

Martin, M.G., 2002. The transparency of experience. *Mind & Language*, 17(4), pp.376-425.

Martin, M.G., 2006. On being alienated. *Perceptual experience*, pp.354-410.

Martin, M.G., 2010. What's in a Look?. *Perceiving the world*, pp.160-225.

McDowell, J., 1982. Criteria, defeasibility, and knowledge. *Studies in the Philosophy of Logic and Knowledge*, 68, pp.455-79.

Miller, A., 2017. Berkeley's puzzle. *Analysis* 77, (4), pp.232-242.

Millikan, R.G., 1984. *Language, thought, and other biological categories: New foundations for realism*. MIT press.

Millikan, R.G., 1996. On swampkinds. In *Mind & Language*.

- Millikan, R.G., 2007. An Input Condition for Teleosemantics?: Reply to Shea (and Godfrey-Smith). *Philosophy and Phenomenological Research*, 75(2), pp.436-455.
- Moore, C.M., Yantis, S. and Vaughan, B., 1998. Object-based visual selection: Evidence from perceptual completion. *Psychological science*, 9(2), pp.104-110.
- Nanay, B., 2010. Perception and imagination: amodal perception as mental imagery. *Philosophical Studies*, 150(2), pp.239-254.
- Neander, K., 2012. Teleological theories of mental content. *Stanford Encyclopaedia of Philosophy*.
- Neander, K., 2013. Toward an Informational Teleosemantics. *Millikan and her critics*. John Wiley & Sons.
- Neisser, U. and Becklen, R., 1975. Selective looking: Attending to visually specified events. *Cognitive psychology*, 7(4), pp.480-494.
- Nida-Rümelin, M., 2007. Transparency of experience and the perceptual model of phenomenal awareness. *Philosophical Perspectives*, 21, pp.429-455.
- Nozick, R., 1981. *Philosophical explanations*. Harvard University Press.
- Price, C., 2000. General-purpose content. *International Studies in the Philosophy of Science*, 14(2), pp.123-133.
- Price, C., 2001. *Functions in mind: A theory of intentional content*. Oxford University Press.
- Pylyshyn, Z., 1989. The role of location indexes in spatial perception: A sketch of the FINST spatial-index model. *Cognition*, 32(1), pp.65-97.

Pylyshyn, Z.W. and Storm, R.W., 1988. Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial vision*, 3(3), pp.179-197.

Quine, W.V.O., 2013. *Word and object*. MIT press.

Quine, W.V.O., 1976. Two dogmas of empiricism. In *Can Theories be Refuted?* (pp. 41-64). Springer, Dordrecht.

Roca-Royes, S., 2011. Conceivability and De Re modal knowledge. *Noûs*, 45(1), pp.22-49.

Rupert, R.D., 1999. The best test theory of extension: First principle (s). *Mind & Language*, 14(3), pp.321-355.

Ryder, D., 2004. SINBAD neurosemantics: A theory of mental representation. *Mind & Language*, 19(2), pp.211-240.

Scholl, B.J., 2001. Objects and attention: The state of the art. *Cognition*, 80(1-2), pp.1-46.

Scholl, B.J., Pylyshyn, Z.W. and Feldman, J., 2001. What is a visual object? Evidence from target merging in multiple object tracking. *Cognition*, 80(1-2), pp.159-177.

Sears, C.R. and Pylyshyn, Z.W., 2000. Multiple object tracking and attentional processing. *Canadian Journal of Experimental Psychology* 54(1), pp.1-14.

Shoemaker, S., 1990. Qualities and Qualia: What's in the Mind?. *Philosophy and Phenomenological Research*, 50, pp.109-131.

Shoemaker, S., 2000. Introspection and phenomenal character. *Philosophical Topics*, 28(2), pp.247-273.

- Shoemaker, S., 2003. Content, character and color. *Philosophical Issues*, 13, pp.253-278.
- Shoemaker, S., 2006. On the ways things appear. *Perceptual experience*, pp.461-480.
- Siegel, S., 2010. *The contents of visual experience*. Oxford University Press.
- Siegel, S., 2016. *The rationality of perception*. Oxford University Press.
- Smith, A.D., 2005. *The problem of perception*. Motilal Banarsidass Publishe.
- Soteriou, M., 2013. *The mind's construction: The ontology of mind and mental action*. Oxford University Press.
- Soteriou, M., 2016. *Disjunctivism*. Routledge.
- Speaks, J., 2015. Is phenomenal character out there in the world?. *Philosophy and Phenomenological Research*, 91(2), pp.465-482.
- Thompson, B., 2010. The spatial content of experience. *Philosophy and Phenomenological Research*, 81(1), pp.146-184.
- Tye, M., 1992. Visual Qualia and Visual Content. *The contents of experience*. Cambridge University Press.
- Tye, M., 1995. *Ten problems of consciousness: A representational theory of the phenomenal mind*. mit Press.
- Tye, M., 2003. Qualia. Stanford Encyclopaedia of Philosophy.
- Tye, M., 2009. *Consciousness revisited: Materialism without phenomenal concepts*. Mit Press.

Tye, M., 2015. Yes, phenomenal character really is out there in the world. *Philosophy and Phenomenological Research*, 91(2), pp.483-488.

Williamson, T., 2007. *The philosophy of philosophy*. John Wiley & Sons.

Yablo, S., 1993. Is conceivability a guide to possibility?. *Philosophy and Phenomenological Research*, 53(1), pp.1-42.

Yantis, S., 1992. Multielement visual tracking: Attention and perceptual organization. *Cognitive psychology*, 24(3), pp.295-340.

Yantis, S. and Hillstrom, A.P., 1994. Stimulus-driven attentional capture: evidence from equiluminant visual objects. *Journal of experimental psychology: Human perception and performance*, 20(1), p.95.