December 2015

Phenomenal Intentionality and the Problem of Cognitive Contact

Christopher A. Young
The University of Western Ontario

Supervisor
Dr. Andrew Botterell
The University of Western Ontario

Joint Supervisor
Dr. Angela Mendelovici
The University of Western Ontario

Graduate Program in Philosophy

A thesis submitted in partial fulfillment of the requirements for the degree in Doctor of Philosophy

© Christopher A. Young 2015

Phenomenal Intentionality and the Problem of Cognitive Contact

by

C.A. Young

Graduate Program in Philosophy

A thesis submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy

The School of Graduate and Postdoctoral Studies
The University of Western Ontario
London, Ontario, Canada

Abstract

Part 1 of the thesis questions the traditional relation model of intentionality. After fixing reference on the target phenomenon, intentionality, and explaining my interest in it, I ask what sorts of things intentionality might be a relation to. I consider ordinary objects, properties, propositions and hybrid views, and conclude all make the intentional relation appear rather mysterious. From there, I move on to examine the relation view's most prominent proponents, the tracking theorists—pointing out some challenges such views face, and concluding that it might be worthwhile looking into alternatives to the relation view.

Part 2 asks whether the newly emerging phenomenal intentionality movement can provide a viable alternative to the relation model of intentionality. After focusing on a specific kind of phenomenal intentionality theory—something I call modificationism—I examine three such accounts. From there I go on to discuss some common complaints/challenges these kinds of views face, and consider how they might be addressed within the modificationist framework.

In Part 3, I address what I call the problem of cognitive contact: how do our contentful mental states manage to make cognitive contact with the ordinary objects (e.g. tables and chairs) that they appear to. The problem is particularly acute for any version of phenomenal intentionality that denies the relation view, and has been given very little attention in the literature. I consider how a modificationist might address this problem, and conclude that though some avenues appear promising, there is nevertheless a great deal of work to be done if modificationism, and phenomenal intentionality theory, is to overcome this problem.

Keywords: Phenomenal Intentionality, Intentionality, Phenomenality, Consciousness, Internalism, Cognitive Contact

Acknowledgements

I am assuming there are unofficial protocols for writing the acknowledgement section of one's thesis. That being the case, I ask your forgiveness for what will undoubtedly (and predictably) be an exercise in not following protocol.

First, I want to express my most sincere thanks to my supervisors, Dr. Angela Mendelovici and Dr. Andrew Botterell. 'Extraordinary' can be used to denote a great many things, and extraordinariness is measured and attributed in just about as many ways. As I want to use it here, it denotes the kind of accommodating, challenging, shrewd and, above all, helpful support that managed to make me a better philosopher. More importantly, it also denotes the attribute of never showing anything but the utmost confidence that I could actually do this; and this even when I was pretty darn sure that I couldn't. As much as anything else, this sense that at least two people believed in my philosophical competence was always a source of that one thing that keeps you going through the dark times: hope. To my supervisors: You are both most extraordinary!

I suppose that by the time one completes his doctoral thesis, his philosophical beginnings are often but a distant memory—like the starting line of a marathon, or one's own driveway after a cross-country journey. Not so with me. I consider my undergraduate days not as merely the distant beginning of this long and arduous trek, but as the formative years that laid the philosophical foundations upon which this project was built—foundations that, for all the weight and abuse they took, never betrayed a fissure. To the man who taught me how to read, write and think, Dr. Peter Campbell: Thank you. *Sine qua non!*

Perhaps unconventionally, yet no less sincerely, I would like to thank Jason Cruz, Jake Kiley, Jordan Burns, Chris Aiken and Rob Ramos. During the darkest times, when I was strung out and stressed to the max, when failure seemed assured, you convinced me to put that next foot in front of the other. You were the peace of mind that kept me going on, got me out of bed, and out that door! And for that, I owe you more than words alone could say.

I seem to remember a story about a philosopher who often became so defeated by the crushing weight of his own revelations that he had to head down to the pub to play cards just to get out of his own head. Earth shattering revelations aside, I too have needed to be brought back to earth by those wise people whose grasp on what is really important in life is eminently firmer than mine will ever be. To Dan, Arlene, Mark, Jenna, Ryan, Noah, Natalie, Sarah, Amy and Lily: Thank-you.

Finally, and most importantly, I want to thank my family. For Matt: I hope the completion of this document is one step closer to making you as proud of me as I am of you, and of having you as a brother. For Sue: For reasons best known to yourself, you always believed in me; and, when I felt I could confide in no one else, you were there. For my parents: Everything good about me is due to you; everything not so good about me is due to my not listening to you. I could never, <u>never</u>, have done this without your love, support and your unfathomable faith that I could. Finally, to my wife Michelle: Though all the forgiveness in the world couldn't pardon the crime of a guy like me convincing a girl like you to marry him, I hope that, small gesture though it is in comparison, the completion of this thesis stands as a tribute to the love, support, and unwavering patience you've shown me.

Dedication

To those who weren't with us too long, Amstel and Duke:

I miss you beyond the count of tears.

This one's for you.

Also, to Dan Redick:

I'm so sorry you couldn't see the end of this project.

Too soon my friend.

Table of Contents

Part 1

Introduction

The focus of Part 1 is the phenomenon of intentionality—the aboutness, or directedness, that some things, systems, or objects display. In particular, I am interested in the species of aboutness displayed by minds/mental states. Traditionally, this phenomenon is thought to be, or involve, a relation between mental states and the items that such states are, or appear to be, about, or directed at. In Part 1, I attempt to raise some concerns about this traditional view.   My plan is as follows: In chapter 1, I attempt to distinguish the kind of intentionality I am interested in from other possible forms of aboutness. I then describe two different conceptions of intentional content, and give a description of content that I think is consistent with both.

If we assume that intentionality is indeed some sort of relation between minds/mental states and distinctly existing things, then two questions arise: 1) What kinds of things does intentionality relate us to? 2) What kind of relation is intentionality?

Chapter 2 addresses the first question—namely, what kinds of things does intentionality relate us to. I consider what I take to be some of the most popular contenders: everyday objects, such as tables, chairs and states of affairs; and abstract objects, such as propositions and properties. I raise concerns for both options.

In chapter 3, I set aside the issue of what kinds of things contents are, assuming a kind of content agnosticism, focusing instead on the most influential attempts at spelling out the nature of the intentional relation. These are the tracking based accounts found in the writings of authors such

as Fodor (1988, 1992), Dretske (1997), and Millikan (2002). With the exception of Millikan's consumer-based theory, tracking theories are almost unanimously causal. I therefore sort these causal theories according to how they attempt to supplement the causal relation so as to allow for misrepresentation. I consider appeals to normal conditions, natural selection, learning and asymmetric dependence—all of which face some concerns.

Part 1 ends with a summary of the results of this foray into what I hereafter will call intentional relationalism. My conclusion will be that though intentional relationalism has been the most widely accepted stance on intentionality, the concerns raised make a search for an alternative view understandable.

1. Intentionality: Fixing Reference

When you read the words on this page, when you wonder whether spring will ever come, when you crave extremely hot chicken wings, you are enjoying mental states that appear to be about, or directed at, things. *This aboutness/directedness is intentionality*. This part of the thesis is about the nature of that phenomenon. In particular, this chapter asks whether intentionality is, or consists in, a relation.

Intentionality is sometimes referred to as mental representation. Neither 'intentionality' nor 'representation' is without its difficulties, nor do the words themselves appear synonymous: While it sounds acceptable to say that smoke *represents* fire, it sounds odd to say that smoke is *about* fire. In my view, the distinction between intentionality and representationality is terminological rather than substantive. All the same, in what follows I will primarily talk about intentionality when trying to describe the aboutness that certain mental states exhibit.

1.1 Original/Derived Intentionality

On the other hand, there is an important distinction between two forms of intentionality, namely intrinsic/original intentionality, and derived intentionality. While this distinction is not uncontroversial,[1] I nonetheless follow a large number of theorists in accepting a real distinction here, and allowing that there are indeed cases of both derived and original

---

[1] See (Dennett, 1987, pp. 288-297), who argues that all intentionality is derived.

intentionality. See, for instance, (Fodor, 1988, pp. 99), (Dretske, 1997, pp. 7-8), (Searle, 1983).

An extremely helpful way of articulating this distinction can be found in (Bourget, 2010): Things have derived intentionality when they "have their contents or representational properties, at least in part in virtue of intentional states distinct from themselves, or relations to such states (Bourget, 2010, pp.33)." Natural language, street signs, and paintings all count as having derived intentionality in this sense: they are about what they are about in virtue of intentional states distinct from themselves—namely our thoughts. [2] On the other hand, some mental states have their intentionality originally. We do not, for instance, *decide* (as we do in the case

---

[2] One problem that emerges for this way of cutting the original/derived distinction is that it implies that intentional mental states that have their contents in virtue of being related to other mental states turn out to be derived. If we take holism—the thesis that something is about what it is about, or means what it means, only relative to the entire (representational) system of which it is a part—seriously, then it looks like the majority of intentionality is dervied. Perhaps this is not such a concern: holism is certainly not a unanimous view. Fodor is a particularly vociferous opponent (see Fodor & Lepore, 1992). However, at the outset, I do not wish to preclude any theory for merely definitional reasons. On the other hand, I think Bourget's way of articulating the distinction accurately distinguishes between different species of intentionality. For holists, I propose the following: *Something has derived intentionality when it has its aboutness, at least in part, in virtue of the aboutness of things disitinct from itself. Something has original intentionality if its intentionality is not derived.* How exactly does this amended formulation of the original/derived distinction avoid the holism problem? Holism can be understood as a thesis about how representational things get to be about *what* they are about, rather than a thesis about how things get their aboutness. One can thus be a holist about *what* exactly thoughts are about—namely they are about *what* they are about in virtue of their relation to the entire representaitonal system— without denying that said thoughts have their *aboutness* originally. To be as transperent as possible: I am not hereby endorsing a view on which the aboutness of some mental state, and *what* that state is about, come cleanly apart. For instance, perhaps words get their *aboutness* in virtue of our deciding *what* they are about. What I am saying is that if you accept some form of holism, you could take the line that there is a real disitnction between a) a mental state's aboutness, and b) *what* it is about, hold that original/derived disitnction concerns only a) and that holism concerns only b), and thereby save the disitnction from holism or vice versa. Holists, and anyone else who thinks intentional mental things *have their contents* in part because of relations they bear to other intentional things, can follow me this far.

of words) that our thoughts have intentionality: they come, as it were, prefurnished.

All that being said, I am willing to concede this formulation of the original/derived distinction, if needs be, and fall back on just citing paradigm cases. Paradigm things that have derived intentionality are words, street signs and pictures. Paradigm things that have original intentionality are mental states. My concern in what follows is original intentionality.

1.2 Introspection and Intentionality

One thing to note about original intentionality is that it is, at least, sometimes observable. We are often able to notice that our thoughts, perceptions, memories, desires, and so on, have the aboutness in question. This is good news, since it allows observation and experience to count among the tools we can use in our account of this phenomenon (see Mendelovici, 2010, pp. 2-7). I am not alone in thinking this (see, for instance, Mendelovici, 2010; Kriegel, 2011).[3]

Of course, the kind of observation appealed to here is a form of introspection, and introspective accuracy is a controversial topic.[4] To forestall worries here, I note the following: When I silently pick a number between one and twenty, there are certain features of my thought that are transparent to me. First, it is clear that my thought, at the very least, purports to have aboutness. Second, I need only look to my own thoughts to observe which number I picked. That is, there is a perfectly clear sense in which *what* number I picked is transparent to me. It is this kind of introspection, that my proposal

---

[3] Kriegel offers an extensive argument that our concept of intentionality is experiential (2011, pp. 3-47).
[4] See (Schitzgebel, 2012), (Smithies & Stoljar, 2012).

will appeal to. What is not transparent to me is why, exactly, I chose the number 19. As I understand it, we are pretty bad at this latter form of introspection: It might, for instance, turn out that a repressed longing to be 19 again caused me to choose an odd number; and this despite the fact that, if asked, I might confabulate some story about a mathematical penchant for primes.

1.3 Intentionality and Content

On the surface, the concept of *intentional content* might seem clear enough. Intentional states display the phenomenon of aboutness, and what they are about is their content. Upon closer inspection, however, things get a bit muddy. If I ask you what the content of your thought is, it seems like I am asking you what your thought is about. Or, if I ask you what you are thinking about, it seems like I am asking you for the content of your thoughts. But how exactly are we to understand this? On one reading, intentional states are about their contents. But this seems like an odd way of speaking: My thought that the grass is green seems to be about the grass, not about this *thing* called a content. I find the following kind of answer unhelpful: "in describing what one's thought is about, namely the grass, one *gives* the content of one's thought." Perhaps less confusing, but still slightly cryptic, one might say that intentional states *have* contents. It is in virtue of your thought's being about the grass that it has the content it does.

What is it then, to *have* content? One answer here is that content arises in virtue of a relation—one that holds between mental states and something else. In virtue of A's bearing a particular spatial relation to B, A is said to be *x distance* from B. But, at least on the present line of thinking, there is not

some *thing* that both A and B are related to—namely this *thing* called a distance.[5]

There is something very intuitive about parts of this view, which goes back as far as Thomas Reid (Reid, 1983, pp. 129-150).[6] More recently, Davidson has argued against what he calls the "meanings as entities" (contents as things) view (Davidson, 1967/2001, p. 20).[7] Something about the idea that contents are not *things*, rings true with me.[8] However, the question of whether contents arise in virtue of being related to things is an additional claim that I will not endorse so early on. For now, I will just note that there is this view of content, and that in what follows, I will try to be sensitive to it.

On the other hand, there are several examples that indicate many theorists have a view of contents on which they are (distinctly existing[9]) things. Here is David Pitt:

> The contents of mental representations are typically taken to be abstract
> objects (properties, relations, propositions, sets, etc.) (Pitt, 2013).

Likewise, Fodor too seems to think contents are *things* in the relevant sense. In his canonical formulation of the representational theory of mind, Fodor writes:

> For any organism O, and for any proposition P, there is a relation R and a
> mental representation MP such that: MP means (expresses the

---

[5] Thank you to Rob Stainton for this useful analogy.
[6] Thank you again to Rob Stainton for pointing me in Reid's direction. It should, however, be noted that Reid's attack on Locke is most plausibly taken to be an attack not just on the view that contents are *things*, but on the representational theory of mind in general.
[7] Davidson is interested in linguistic meaning, not (underived) intentional contents.
[8] I discuss this in much further detail in Part 2.
[9] This is Mendelovici's terminology (see her MS, ch. 8)

proposition that) P; and O believes that P iff O bears R to MP (Fodor, 1992, p. 16).

Diagrammatically, Fodor's view can be represented thus:
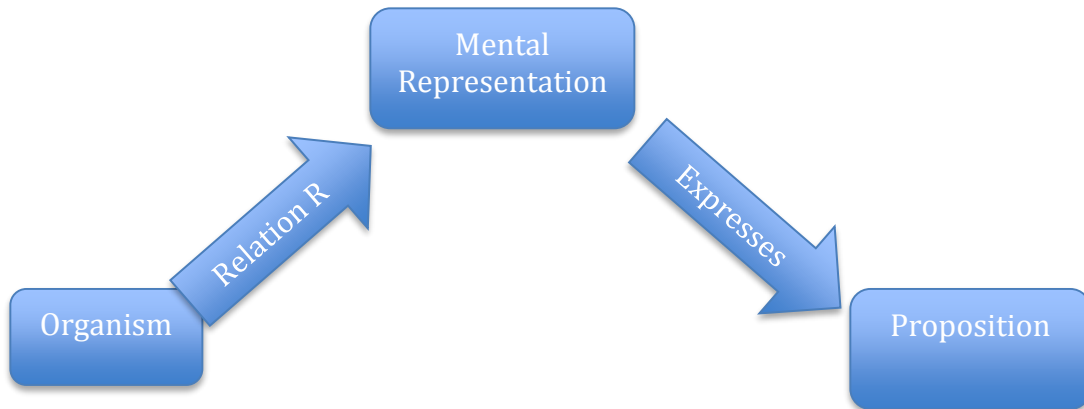
It seems pretty clear what the content is supposed to be here, and that the mental representation bears the expresses relation to *it*. In this case, it looks like the content of the mental representation is the proposition. And as Fodor is a realist about abstract entities such as propositions,[10] it would seem that he takes intentionality (mental representation) to be a relation to independently existing things that play the role of contents.

This is important for what follows, since I will be questioning whether or not intentionality is a relation. The way I propose to go about questioning this relational view of intentionality is by first examining what kinds of things intentionality relates us to. And I follow several theorists (Mendelovici, 2010; Pitt, 2009; Kriegel, 2011a,b) in understanding this examination as an examination of what kinds of things can play the role of contents. If the above examples are any indication, I think this way of describing things is not too

---

[10] See (Fodor, 1992, pp. 132 n. 6).

far off base. However, if the reader dislikes this particular way of speaking, then s/he is free to understand what follows thus: The relational view of intentionality has it that intentionality is a relation. Relations have relata. One of the relata in the intentional relation is our mental states. Chapter 2 is concerned with what exactly the other relatum is.

All that said, I have still not provided even a tentative, reference fixing, account of content. Along with (Kriegel, 2011, pp. 151-152), and (Pitt, 2009), I start by noting that among other things, we often invoke or appeal to content in order to distinguish one intentional state from another. One of the differences between imagining a Labrador and a Pug is a difference in content. We might thus approach content by saying that *intentional content* is one way that an intentional state differs from another, or something that distinguishes one intentional state from another. Though rather thin, this initial way of approaching content is silent about the nature and mechanics of intentionality. On this view, intentionality could turn out to be essentially a relation between mental states and distinctly existing items (e.g. ordinary objects, propositions, property instances) that serve as contents. On such a view the difference between imagining a Labrador and a Pug would be that in the former case, one's intentional state is related to the, e.g., the property of being a Labrador, and in the latter, the property of being a Pug. On the other hand, taking content to be one way intentional states can differ is also compatible with the view that contents are not distinctly existing things. For instance, quickly and slowly are ways that runnings can differ, but quickly is not a distinctly existing thing.

1.4 A brief Recap

So far, I have explained that my target phenomenon is a species of intentionality, or aboutness—namely, original intentionality. Something has original intentionality if its intentionality does not derive from other intentional things distinct from itself. Words, street signs and pictures fail to have original intentionality, and therefore fall outside the focus of this thesis. I have also noted two possible ways of understanding the concept of intentional content, and provided the reader with two possible ways of understanding the project of investigating what sorts of things intentionality relates us to. Finally, I proposed an encompassing conception of intentional content.

2: The Objects of Intentionality

Any view that takes intentionality to be a relation between minds and distinctly existing things is a *relation view*. Different relation views disagree about what items play the role of contents; and, depending on the particulars of the view, there may be a second relation between contents and the things to which they refer.[11] In this chapter, I look at the various options for specifying the objects to which we are related via intentionality.

2.1 Intentionality and the DR-relation view

Pre-theoretically, intentionality looks like a phenomenon that relates us to items and states of affairs in the world beyond our skins—most often everyday items.[12] When I say of someone that she often thinks about her Labrador, Buddy, it looks like I am saying that she often bears a relation to Buddy. Pre-theoretically then, it looks like when we say X is about Y (where X is some intentional mental state), the 'is about' locution names a relation between the mental state X and some ordinary object, Y, such as a cup of coffee or a Labrador.

Because this pre-theoretical view enjoys a strong kinship with direct realist theories of perception—indeed it is basically direct realism writ large enough to encompass all forms of intentionality (perceptual, cognitive, etc.)—I will call it the DR-relation view. **The DR-relation view** takes intentionality to be a relation between mental states (or subjects) and everyday items.

---

[11] I am speaking, of course, about views that take contents to be something like abstract, perhaps universal, properties.

[12] 'Items' should be read to mean not only artefacts, but also states of affairs. Thus, Barak Obama, a coffee cup, and the dog being on the couch all count as items.

Several challenges emerge for what I am calling the DR-relation view:

1) I can think about something that does not exist. Often referred to as the problem of intentional inexistence (Brentano, 1887/1995), this worry has probably been around as long as philosophy. More recently, the worry has been raised by (Mendelovici, 2010; Kriegel, 2007, 2008, 2011; Crane, 2001).

2) I can think about something under one aspect, or description but not another. Lois Lane can hope to see Superman tonight, without thereby hoping to see Clark Kent (Mendelovici, MS)

3) I can think about indistinct things. (see Mendelovici, MS; Kriegel 2008). I can desire a sloop without desiring any particular sloop (Quine W. , 1966, pp. 185-186), I can think of a man without thinking of a man of any particular height (Anscombe, 2002, p. 58), and I visualize a tiger without visualizing a tiger with a particular number of stripes (Dennett, 1969/1986, pp. 136-137).[13] But I cannot feed an Anscombian man to a Dennettian tiger while sailing on a Quinean sloop.

In what follows, I focus on 1) and 2)—though more time will be spent on 1) since I take it to be the most important.

2.1.1 Intentional Inexistence and the DR-relation View

Above, I said that the DR-relation view takes intentionality to be a relation to ordinary/everyday objects. Among the ordinary objects I have in mind are

---

[13] It should be noted here that Quine, Anscombe and Dennett do not raise the above examples in the context in which I am putting them to use. I use these examples because they are relatively famous.

tables and chairs, people, and so on. So, the DR-relation view holds that for me to think about my kitchen table is for me to stand in a relation to my table. Of course, I doubt very much whether anyone thinks that all intentionality is a relation to ordinary objects. This would imply, for instance, that I cannot have thoughts about numbers, which are plausibly taken to be abstract objects. However, the concerns that I shall discuss below apply equally to more restricted DR-relation views—ones that take only certain forms of intentionality to be relations to ordinary objects. For instance, one might hold that singular thoughts—thoughts about J.K. Rowling, thoughts about your kitchen table—have these objects as constituents. Or, someone might hold that all those thoughts that purport to be about ordinary objects are relations to those ordinary objects.

The problem of intentional inexistence is that I can entertain thoughts about unicorns, Bigfoot, Pegasus, thestrals, Santa Claus, and so on. If thinking about something involves being related to it, how can I think about something that does not exist? In what follows I will focus almost exclusively on the example of unicorns. However, for a DR-relation view whose scope covers only singular thoughts, just replace every instance of 'unicorn' with 'Pegasus'.

A useful way to think about this problem is in terms of an inconsistent triad of sentences (see Kriegel, 2007, pp. 307-308):

1) I can think about things that do not exist.
2) I cannot be related to something that does not exist.
3) Thinking about something = being related to it.

The question is how to resolve the inconsistency of the triad. To deny 3) would be to abandon the DR-relation view, or any other relation view for that matter. The options for the DR-relation view are thus to deny 1) or 2).

A denial of 1) could take several forms, most of which are vetted and ultimately rejected in (Kriegel, 2007, pp. 310-311). Considered first is the view that when we take ourselves to be thinking of unicorns, we are actually not thinking at all (Kriegel, 2007, p. 310). Though Kriegel finds this view highly implausible, one might take this line because of other theoretical commitments. For instance, someone who accepts a purely referentialist theory of content (see Fodor, 2008), and who accepts semantic externalism about natural kind concepts (see Putnam, 1975), might argue that 'unicorn' is a natural kind concept whose content is therefore its external world referent.[14] Since no such referent exists, 'unicorn' is contentless. Unicorn thoughts, on such a view, are not bonafide thoughts at all. Whether or not such a view is plausible, Kriegel insists that it simply pushes the problem back:

> "…the problem will resurface for the activity of *seeming to oneself to be thinking of something* (in the relevant sense of "seeming," where an intentional mental state is actually attributed). That is, we can devise a new inconsistent triad: one can seem-to-think of non-existents; one cannot bear relations to non-existent; yet seeming-to-think of something involves (constitutively) a relation to it." (Kriegel, 2007, p. 310)

One way of denying 1) that Kriegel does not discuss is to argue that when we seem to ourselves to be thinking about, for instance, unicorns, we are actually

---

[14] However, see (Kripke, 1972, pp. 156-157) for an argument that 'unicorn' cannot be a natural kind concept.

entertaining a composite thought that bears several relations to distinct things. In the unicorn case, the things we are related to are horses, and horned animals.

This composite line has two parts: a) A negative thesis about what we take ourselves to be thinking of: When we take ourselves to be thinking of unicorns, we are not thinking of what we take ourselves to be thinking of. And b), a positive thesis about what we are in fact thinking in such cases: When we take ourselves to be thinking of unicorns, we are actually thinking of the composite horses + horned animals. I have concerns about both a) and b).

A) Appears to imply that there are some thoughts—thoughts that seem to us to be about unicorns—that are never about what we take them to be about. As Kriegel points out, any approach that makes such a claim seems to imply a second-order error theory about intentional states (Kriegel, 2007, p. 310). While I follow Kriegel in finding such a view highly counter–intuitive, it is an option. To avoid the counter-intuitiveness, one would have to devise some principle whereby we could distinguish cases of accurately taking oneself to be thinking of x, from inaccurate cases, such as taking oneself to be thinking of y, where y does not exist.

Perhaps this could be done, but the relevant principle would have to distinguish such cases on grounds other than the existence or lack thereof of the entities in question. That is, it would seem circular to say that when you take yourself to be thinking of unicorns, you have to be mistaken (about what you are thinking about) because unicorns do not exist. Recall that we are here considering how someone who holds the DR-relation view might deal with the inconsistent triad of sentences above. In particular, we are

considering whether someone might reasonably deny that we can think of things that do not exist. The present proposal is that when we take ourselves to be thinking of unicorns, we are mistaken. Unless someone keen on such a proposal is ready to bite the bullet and accept a full-blown second-order error theory about intentional states—a theory that says we are often/always wrong about what we take ourselves to be thinking about (viz. what we take ourselves to be thinking is rarely, if ever, what we are in fact thinking)—then some principle whereby we might distinguish the bad cases from the good ones seems required. But the relevant principle cannot be simply that the bad cases are those where what one takes oneself to be thinking about does not actually exist. Otherwise, the explanation is circular. Again, divining such a principle might be possible, but I think the burden of evidence is on those who deny that we can think about things that do not exist.

2.1.2 A Word on Externalism

On the other hand, one of the central lessons (or consequences) of externalism is that what we are in fact thinking of is determined, in part, by external factors—factors that can be quite outside a subject's epistemic reach, yet no less determinative of what s/he is thinking. Hence, it should come as no surprise that we can and often are mistaken about what we take ourselves to be thinking of. I cannot say everything I want to about externalism at this point. However, I will note that a great deal of ink has been spilled, by a great many important thinkers, trying to reconcile some form of externalism with self-knowledge (see Davidson, 1996; Burge, 1996; McKinsey, 1996; Bilgrami, 1996). Indeed, that in his seminal paper (Putnam, 1975) Putnam himself opted for a narrow/wide bifurcation of content has seemed to some as an attempt to allow for self-knowledge (see Bilgrami,

1996). This alone suggests that throwing out self-knowledge with the internalist bathwater is a consequence that even the most ardent externalists (e.g. Putnam, Burge) are wary of. And to incur such a loss just to account for thoughts about things that do not exist seems worse.

An interesting worry for the externalist in this context can be found in (Boghossian, 1998). In order to bring Boghossian's point to bear on the present concern, I adapt his insights in the following brief thought experiment: Ed and Ted are internal duplicate unicorn thinkers, living on earth and twin earth respectively. The difference between earth and twin earth is that the latter contains unicorns. According to the present line, when Ed takes himself to be thinking of unicorns, he is actually entertaining a composite thought. Let us assume that having a composite thought is a matter of tokening a molecular concept: 'horses with horns'. On the other hand, when Ted takes himself to be thinking of unicorns, he is tokening the atomic concept 'unicorn'. The problem is that the external factors that are supposed to determine content—the relations or lack thereof between Ed/Ted and unicorns—are not *just* having an effect on content; they are having an effect on syntactic *form*. That is, Ed's unicorn concept is molecular, and Ted's is atomic. And this appears to be a rather serious consequence for any view that takes syntactic form to be a matter of internal constitution: It is part of any twin earth thought experiment to hold internal states constant. Therefore, syntactic form is not (wholly) internal in the same sense that meaning and content are not. It might be objected that I am incorrectly assuming 'unicorn' to be atomic on Twin Earth, and that this argument depends on that assumption. On the contrary, take any earthly atomic concept, 'A', and have it be tokened by Ed. Now assume a Twin earth where whatever normally causes the tokening of 'A' on Earth does not exist. The

composite line holds that because there are no A-s on Twin Earth, when Ted tokens 'A' he is tokening a composite/molecular concept.

In a similar vein, Segal (2000, pp. 54) raises concerns about how an externalist might spell out the extension conditions for the concept 'unicorn'. There are two options according to Segal:

> "The first is that such concepts are modally empty: they do not apply to anything any possible world. The other is that they are motley concepts, applying, roughly speaking, to anything satisfying the core descriptions associated with them" (Segal, A Slim Book About Narrow Content, 2000, p. 54).

While we are considering the latter possibility—namely that unicorn thoughts are actually composite thoughts about horses and horns—it is worthwhile noting that the former option entails that unicorn thoughts are on par with thoughts about round squares and Penrose triangles. No extension conditions = modally empty = impossible. I suppose this is a live option, but further argument would be required. Returning to the second option—the motley, or composite, concept view—I begin by noting that fixing extension conditions by way of core descriptions seems like it would be an uncomfortable prospect for an externalist. The more pressing point, according to Segal, is that any extension fixing core description will pick out the same things across worlds, pace externalism (see Segal, 2000).

Again, the point here is not to mount a full scaled examination of the merits of externalism. Rather, it is to discover whether externalism can quickly and decisively come to the aid of the DR-relation theorist who seeks to account for our seeming ability to think about things that do not exist by denying that

we can think about such things. The proposal is that when we think about unicorns, we are actually thinking a composite thought about horses and horns. The initial worry was that this involves unicorn thinkers in a kind of second order error about what they take themselves to be thinking about. The appeal to externalism is supposed to temper this consequence because externalism—a widely excepted view—implies that we can and often are mistaken (or otherwise lacking complete knowledge) about what we take ourselves to be thinking about. The first point I made was that the rejection of self-knowledge is not taken to be unanimously unproblematic, even by the most ardent externalists. It is clear from the works of (Davidson, Knowing One's Own Mind, 1996), (Burge, Individualism and Self-Knowledge, 1996), (McKinsey, 1996), (Bilgrami, 1996), (Kriegel, 2007) and (Putnam, 1975) that self-knowledge is not something to be discarded lightly. The second point was that authors like Segal (2000) and Boghossian (1998) have provided some powerful reasons to think that externalism has problems accounting for empty concepts. Hence, an appeal to externalism in the present context might not be appropriate. At the very least, externalism does not supply a quick and decisive way of denying that we can think about things that do not exist.

To come back to the issue at hand, the view we are examining is one on which intentionality is a relation our mental states bear to ordinary objects such as tables and chairs—a view I have called the DR-relation view. From this, and the fact that we can think about things that do not exist, we have the following inconsistent set of sentences:

1) I can think of things that do not exist.
2) I cannot be related to something that does not exist.
3) Thinking about something is a way of being related to it.

I have raised some concerns about different ways of denying 1); but might someone deny 2)? There have been several theorists, both historical and recent, who have sought to do so (see Parsons, 1975; Meinong, 1960); but none take the view that intentionality is a relation to ordinary objects: The objects in question are construed as having being, but not existence (Meinong, 1960), which makes them extraordinary.

On the other hand, it is possible to argue that I can be related to an ordinary object that does not exist in this world, but does in another. The thinking here would be to accept a strong modal realism along the lines of David Lewis (1986); allow that things like unicorns, and bigfoot, number among the ordinary objects in some possible world; and construe the intentional relation to hold between you in this world—where unicorns do not exist— and said things in the world in which they do exist. Whether or not this view is ultimately plausible, I leave to the reader. However, it should be noted that this makes the intentional relation look rather fantastical—able to stretch between worlds. Certainly it makes the intentional relation different from some other relations such as causation.

2.1.3 The DR-relation View and Thinking About Things Under an Aspect or Description

Relations between ordinary objects are not sensitive to description in the way intentionality is. I cannot kick Superman without thereby kicking Clark Kent. Likewise, I cannot be taller than Superman without thereby being taller than Clark Kent. I can, however, hope to see Superman without hoping to see Clark Kent. There are several arguments in the literature that raise this point (see Mendelovici, MS; Kriegel, 2011).

Consider the following: Assume that I am unaware that the morning star is the evening star, so that I can think about the latter without thinking about the former (for instance, I can hope to see the one without hoping to see the other). For any other (non-intentional) relation I bear to the evening star I also bear it to the morning star. I cannot point to the evening star without pointing at the morning star; I cannot shoot at the evening star tonight without shooting at the morning star; I cannot be less massive than the evening star without being less massive than the morning star, and so on. If thinking about the evening star is a matter of bearing the thinking-about relation to the evening star, then the thinking-about relation is unlike any other non-intentional relation: It relates me to the ordinary object (the evening/morning star) under one description/aspect, but not another (see Mendelovici MS, Kriegel 2011, pp. 127 – 132).

2.1.4 Summary

There is undoubtedly more one could say about the DR-relation view and intentional inexistence, but I think I have raised some fair concerns for the view. At the very least, the DR-relation view must choose to: 1) Deny that intentionality is like any other relation that holds between everyday objects. Call this the mysterious relation view. 2) Allow that in some cases (e.g. non-existent objects, objects under a particular description) intentionality does not relate us to ordinary objects, but rather to abstract properties, universals, propositions, or some other non-everyday item. Presumably, the mysterious relation view—being a species of what Kriegel (2007, pp. 311) calls "intentionality exceptionalism"—is something to be avoided: It sounds odd to say that certain facts hold of every relation between ordinary objects except

for one (see Kriegel, 2007, pp. 311).[15] However, on 2) the DR-relation view is no longer, as it were, pure: In some cases, intentionality relates us to something other than ordinary objects. Call this second line the Impure DR-relation view. I will address the Impure DR-relation view below, but I will first examine whether abstracta are good candidates for the things intentionality relates us to.

2.2 Abstracta as Intentional Contents

There are, of course, alternatives to the DR-relation view. Indeed, advocates of DR-relation views are in the minority. More common is the view that intentionality is a relation between mental states and abstract objects such as properties or propositions. Before examining the merits of views that take intentionality to be a relation to such things, it should be noted that the literature concerning properties, propositions, abstract objects, and so on, is vast. For instance, it has been held that the particular/property (universal) distinction might not be as clear as it is often thought (see Ramsey, 1997). This would make any examination of the purported relation between intentional states and the properties they are about that presupposes a clear distinction between particulars and properties wrong-footed from the get go. There are also views of properties on which properties are classes or sets (Lewis, 1997, p. 190). On such view, the pre-theoretical sense that properties are things that particulars have—or parts of particulars—is turned on its head:

"Far from the property being part of the [particular], it is closer to the truth to say that the [particular] is part of the property. But the precise

---

[15] I will consider the possibility that intentionality is a unique relation below.

22

truth, rather, is that the [particular] is a member of the property (Lewis, 1997, p. 190)."

Propositions are sometimes thought to be particular kinds of properties—namely, "those that are instantiated only by entire possible worlds (Lewis, 1997, p. 175)." It is also sometimes held that properties are not all that abstract—at least not as abstract as say Platonic universals. Trope theory, for instance, can be seen as denying the kind of full out abstractness to properties that Platonism ascribes them. Though theorists like (Campbell, 1997), (Williams, 1997) call tropes "abstract particulars," it is clear that the degree of abstractness such entities are said to have is significantly less than universal properties. Universals of the Platonic kind exist outside of time and space; tropes do not. What makes tropes abstract—at least on views like Campbell's (1997, pp. 136-137), is that they can contemporaneously occupy the same space as other tropes: My Pyrenees' whiteness is present in the same place as my Pyrenees' furriness.

Even if we could settle on one plausible view of the nature of abstract entities—that they are universals, tropes, sets, or something else, other problems arise. Numbers are often considered abstract, as are propositions, the average taxpayer, the 2004 Kia Sorento, tallness, redness, and so on. But to group all of these things into one category—abstracta—seems problematic. For instance, in reading the owner's manual, I discover that the 2004 Kia Sorento comes equipped with four airbags, but no proposition can be equipped with such things. Likewise, the number four is the product of primes, but the average taxpayer is not. More importantly, abstracta are not members of one causal kind; some are thought to have obvious causal influences in the realm of concreta; others do (may) not. Since causal powers are one way to type objects, it seems that abstracta do not form a single type.

With such disparate views on the nature of abstracta, the only hope of providing an adequate examination of the suitability of abstracta to be that to which we are related via intentionality, is to stay focused on precisely that. Note also, that the suitability, or lack thereof, of abstracta to be one of the relata of the intentional relation does not entail anything about the truth of any given theory of properties, universals, tropes, and so on. It could, for instance, turn out that though tropes are poor candidates for being that to which we are related via intentionality, the trope theory of properties is nevertheless true. Only theories that invoke the relevant abstracta in order to account for the intentional relation would be impugned/vindicated by what follows.

What I propose to do in this section is examine three views of abstract entities: as Platonic properties, as Aristotelian properties and as propositions. Before going on, I should explain why I choose not to address the view that takes the relevant abstracta to be tropes: I will be raising concerns about the suitability of properties that depend on instantiation (a-properties) to be one of the relata in the intentional relation. The concerns I raise there apply equally to tropes.

2.2.1 Intentionality as a Relation to the Forms

Following (Armstrong, 1997, pp. 108) and (Kriegel, 2011b, pp. 252), we can call properties that exist outside of space and time, *ante rem,* and properties that depend on instantiation, properties *in re.* The former are what Plato calls the forms, and have more recently been invoked by philosophers such as Russell (Russell, The World of Universals, 1967/1997) to account for similarity (or dissimilarity) between concrete particulars.

Several metaphysical/ontological challenges arise for those who posit what I will hereafter call Platonic properties—p-properties for short. Quine is notoriously anti-*ante rem*, so to speak (Quine W.V.O., 1953), as is (Devitt, 1997). Both argue against ontological commitment to such things. Indeed, from a physicalist point of view, it would seem that p-properties cannot be tolerated. From (Shoemaker, 1997), one might construct an indirect argument against p-properties. Shoemaker notes that observing something is one way to be causally affected by it (Shoemaker, 1997). If we assume that some properties are observable—not such a stretch given that observed similarities/differences is one of the chief purposes of invoking properties—then some properties are causal. However, if properties are p-properties and exist outside of space and time, then it is difficult to see how they could exert a causal influence in the realm of spatiotemporal concreta. But I digress; what is important here is whether or not p-properties can be one of the relata in the intentional relation.

To be sure, the problem of intentional inexistence disappears with the positing of p-properties: When I think about unicorns, I bear the thinking-about relation to the abstract p-property of unicorness, or perhaps the abstract properties of horseness and hornedness, or some such. Likewise, when I think about the morning star, I bear the thinking-about relation to some p-properties; and when I think about the evening star, I bear the relation to others.

So far, so good. However, a number of issues arise here. First—stemming from the argument we adduced on Shoemaker's behalf just above—the causal impotency of p-properties in the spatiotemporal realm would appear

to preclude providing a causal theory of intentionality (Kriegel, 2011b).[16] Second, one might reasonably ask what kind of relation intentionality is such that it can traverse the divide between the realm of spatiotemporal concreta, and Plato's realm of the forms, a fantastical one to be sure.[17] Of course, Plato's own story is that before we take corporeal form (before we are born), we are in direct contact with this posited realm, and that our thoughts are recollections of the p-properties we encountered there. However, I am not sure whether this would satisfy many current theorists. For one thing, it is patently dualistic.[18] For another, it ties the intentional relation essentially to memory. All thoughts turn out to be species of recollection. Russell's account of how exactly we become acquainted with p-properties is by abstraction:

> "When we see a white patch, we are acquainted, in the first instance, with the particular patch; but by seeing many white patches, we easily learn to abstract the whiteness which they all have in common, and in learning to do this we are learning to be acquainted with [the universal] whiteness (Russell, The World of Universals, 1967/1997, p. 51)."

Of course, Russell was not focused on the same issue I am; but nevertheless, it is unclear whether this kind of answer is adequate to account for how intentionality relates mental states to p-properties. For one thing, this kind of answer seems to require contact with instances of the relevant property. And there again, we seem to run into difficulties with properties that are

---

[16] I do not hold a causal theory, but for those who do, p-properties might not be suitable candidates for intentional contents.

[17] Mendelovici (2010; MS) raises similar complaints. Her arguments focus on the mysteriousness of saying when two things from different ontological categories belong together.

[18] Being dualistic does not necessarily condemn a view. It is just that contemporary philosophical orthodoxy is not dualistic. Hence a theory that implies dualism would undoubtedly enjoy a less than enthusiastic reception.

never instantiated, and so are not the kinds of things that we may abstract from.

Third, whether I am thinking about unicorns, dogs, tables or chairs, the intuitive answer as to what kind of thing I am thinking about is that I am thinking about ordinary concrete objects, not abstracta. As Kriegel notes, phenomenologically, these things present themselves in thought as ordinary concrete things (Kriegel, 2007, p. 310). So this view too appears to imply fairly widespread error about what we take ourselves to be thinking about.

Fourth, Kriegel explains that on such a view, there emerges a "veil of abstracta" over the realm of ordinary concrete objects (Kriegel, 2007). His argument focuses on perception and proceeds in several steps. First, he distinguishes between two models of the relation that holds between perceptual experiences and perceptual beliefs: the inference model, and the endorsement model. On the inference model, we infer beliefs about the objects of perception from our perceptual experiences. On the endorsement model, "some perceptual beliefs are justified simply by taking at face value one's current perceptual experience" (Kriegel, 2011b, p. 247). According to Kriegel, for most perceptually based beliefs,[19] the correct model is the endorsement model (Kriegel, 2011b, p. 247).

While Kriegel does not offer much in the way of argument for favouring the endorsement model, reasons are not hard to divine. First, inference requires a level of intellectual sophistication that certainly eludes animals, small children, and probably some adults. Hence, the inference model implies that the way animals, small children and some adults form perceptually based beliefs is different in kind from the way the rest of us do. Second, inference is

---

[19] Kriegel recognizes certain instances where the inference model is probably better equipped: e.g. scientific beliefs about sub-atomic particles, etc.

more temporally expensive than endorsement, and so is rather a hindrance both evolutionarily, and in our day-to-day lives: imagine how the inference model would account for driving a car. For most of our perceptually based beliefs, such beliefs are acquired by endorsement of our perceptual experiences.

The second step is the veil thesis:

> "There is a class of entities $X_1,…,X_n$, such that for any perceptual experience E and some perceptual belief B, (i) E does not bear an epistemically and subjectively relevant intentional relation to a member of $X_1,…,X_n$, and (ii) B does bear an intentional relation to a member of $X_1,…,X_n$ (Kriegel, 2011b, p. 249)."

Kriegel adds that a relation is epistemically irrelevant "if the experience would justify the same beliefs even if it did not bear it," and subjectively irrelevant "if the subject could not tell the experience apart from another, otherwise similar experience" (Kriegel, 2011b, p. 249). The problem with the veil thesis is that it denies that the objects about which we have perceptually based beliefs—the ordinary objects that furnish our surroundings—are the same objects that we perceive. It therefore mandates the inference model: we infer our beliefs about the concrete objects around us from the objects we *perceive*, where the former is different from the latter. This, as Kriegel notes, is one of the central concerns for the sense data theory (Kriegel, 2011b, p. 248).

But notice that the veil of abstracta only arises when the property theory of perceptual intentionality is paired with the view that perceptually based beliefs are not about the relevant properties. But why would someone who

seeks to account for intentionality in terms of relations to p-properties not think that beliefs are such relations? That is, why would someone construe the perception relation as obtaining between perceptions and p-properties, but the belief relation as holding between beliefs and the ordinary concrete objects that furnish our surroundings? Taking one's beliefs about tables and chairs to actually be beliefs about p-properties is a live option, and one that would undermine Kriegel's argument. The so-called veil emerges because one's perceptions are of p-properties and one's perceptually based beliefs are not; hence the justification for one's perceptually based beliefs must be by inference, not by endorsement; and that is the wrong model. But the veil would not emerge for someone who thinks perceptually based beliefs are about p-properties.

Of course, this may be a moot point, since I doubt anyone thinks that one's beliefs about the ordinary objects in her surroundings are actually about p-properties. But that alone would count as yet another concern for the p-property theory of intentionality. At any rate, if you assume that your perceptually based beliefs are beliefs about the ordinary concrete objects in your surroundings, and that perception is a relation to p-properties, then the veil of abstracta does arise.

To summarize, the view we have been considering is that intentionality is relation to p-properties—universals that exist outside of space and time. Setting aside whatever metaphysical/ontological issues arise by positing such things, there are several reasons why p-properties are not good candidates for being that to which we are related by intentionality. First, the causal impotency of non-spatiotemporal entities (e.g. p-properties) would seem to preclude giving a causal account of intentionality. Second, this view makes the intentional relation look a bit fantastical—able to cross between

29

the realm of ordinary spatiotemporal concreta and the realm of the forms. Third, the view conflicts with the phenomenality of intentional episodes about ordinary concrete objects: such things present themselves in experience as ordinary spatiotemporal concreta, not abstracta. Hence, this view too threatens introspective knowledge—probably not something to be given up just to allow for our representational contact with the realm of the forms. Fourth, if we assume our perceptually based beliefs are beliefs about the ordinary concrete objects in our surroundings, but that perception is a relation to p-properties, then Kriegel's veil of abstracta arises. On the other hand, the alternative to erecting the veil—denying that our perceptually based beliefs are about the concrete ordinary objects we think they are, but are instead about the very p-properties the perceptions are about—seems worse. It seems fair to say that p-properties are not unproblematic candidates for one of the relata in the intentional relation.

2.2.2 Intentionality and properties *In Re*

Properties *In Re* do not exist outside of space and time, but instead are dependent on the concrete objects that instantiate them. In this sense, they are akin to Aristotelian universals; and I will thus henceforth call them a-properties. The metaphysical/ontological virtues of a-properties over p-properties seem obvious. Metaphysically, because a-properties exist within the spatiotemporal realm, their causal efficacy seems less problematic than that of p-properties. Ontologically, we need not posit a distinct realm in which to put a-properties; they are already somewhere: here. Among the most well known proponents of a-properties over p-properties is Armstrong (Armstrong, 1997, 1997b, 1999). The question is whether a-properties can serve as one of the relata in the intentional relation. Certainly, the fact that a-

properties exist *in* the very things that instantiate them appears to favour them as plausible intentional relata.

One concern about the a-property theory of intentionality concerns the relation of a-properties to the concrete things that instantiate them. One might reasonably ask whether, if stripped of all its a-properties, the concrete particular would remain? Armstrong calls such a stripped down particular a "blob" (Armstrong, 1999, p. 199), and argues against the "blob" view—which he claims arises for the p-property view—by having "a thing's properties as constituents of the thing" (Armstrong, 1999, p. 199). This raises the question of whether the concrete objects whose constituents are a-properties are sets of a-properties, or sums of a-properties (see Kriegel, 2011b, pp. 257). I realize that there are some tremendously difficult metaphysical questions concerning this issue; but again, my focus is on the view that takes a-properties to be one of the relata in the intentional relation. Admittedly, on either the set or sum view, the intentional relation to the relevant a-properties gets us to the right object—be it a set or a sum of a-properties. However, I agree with Kriegel that the a-property theory should opt for sums over sets:

> "Given that sets are non-spatial entities,…if…beliefs were intentionally related to sets of i-universals [a-properties], they would be intentionally related to non-spatial entities" (Kriegel, 2011b, p. 257).[20]

On the other hand, both the set and sum views appear vulnerable to extensions of the arguments for content determinacy found in (Horgan & Graham, 2012). I will not rehearse their arguments here, but will note that if

---

[20] In note 37 (Kriegel, 2011b, pp. 265), Kriegel cites (Van Cleve, 1985). Van Cleve's article is both forceful and illuminating in its vetting and rejecting of both the set and sum views of particulars.

certain forms of intentionality can determine that a rabbit-thought is not about the set of undetached rabbit parts, then perhaps it can also determine that rabbit thoughts are not about sets or sums of rabbit a-properties.

However, my main concern about the a-property theory is that, unlike the p-property theory, it reintroduces the problem of intentional inexistence—not for ordinary concrete objects, but for a-properties. A-properties depend on instantiation in a way that p-properties do not. On the a-property view, there can be no uninstantiated properties. However, I can think about unicorns even though the property of unicorness is not instantiated, and that appears to tell against the view that intentionality is a relation to a-properties. I have already discussed the plausibility of explaining away apparent unicorn-thoughts in terms of horses and horns, and will not rehearse my concerns again.

Perhaps, however, you are unsatisfied about properties such as unicorness counting as genuine properties. That is, perhaps you think that as unicorness is not a genuine property, it is not an appropriate example of the shortcomings of the a-property theory of intentionality. While I disagree, the concern can be raised about other, more conventional properties. Take colors, for instance. Colors seem to be properties par excellence. On some views of color (see Hardin, 1988; Mendelovici 2010; Boghossian & Velleman, 1997), color properties are never instantiated and so stand as counter-examples to the view that intentionality is a relation to a-properties (assuming that we have intentional episodes of color). However, the example need not rest on the truth of color irrealism. Kriegel (2011b) points to an article by Churchland, who discusses experiences of certain "chimerical colors" that are not just uninstantiated, but are nomologically impossible (Churchland, 2005). The crucial point is that a-properties depend on

instantiation; there are no unistantiated a-properties. But since I can have thoughts, beliefs, perceptual experiences, and so on, of uninstantiated properties (e.g. chimerical colors), a-properties seem ill-equipped to be the relata of these purported intentional relations.

In summary, there are several worries about the suitability of a-properties to serve as relata in the intentional relation. Positing a-properties involves the theorist in taking a stance on the nature of the relation between the a-properties and the concrete things that have them. Neither the set, nor the sum view of this relation is ideal: The former turns ordinary concrete objects into abstracta by identifying them with sets. And both appear to be phenomenologically inadequate: When I think about a rabbit, I take myself to be thinking about an ordinary concrete object, not a set or a sum of abstract (though instantiated) properties. Finally, and most importantly, the fact that the a-properties depend on instantiation raises, once again, the problem of intentional inexistence.

2.2.3 Intentionality and Propositions[21]

In what follows in this section on propositions, I will be assuming that there is an intuitive sense in which propositions are distinct from properties, pace Lewis (1997, p. 175).[22] Once again, my goal here is to examine the suitability

---

[21] I am here concerned with propositions as the things we (our mental states) are related to via intentionality. Propositions have been assigned a variety of roles outside of the philosophy of mind, but those are not my primary concern here.

[22] If propositions were not distinct from properties—if propositions were, as Lewis would have it, properties "instantiated only by entire possible worlds," (Lewis, 1986, p. 53)—then on the view under consideration intentionality would be a relation to such properties. But, as Lewis says, propositions, on such a view, would just be the set of possible worlds where the proposition is true. And that means that intentionality is a relation to sets of possible worlds, which, according to Lewis, looks to be false: It precludes certain *de se* thoughts on account of failing to discriminate between individuals at the same world (Lewis, 1986, pp. 55).   So, though propositions are

of propositions to be one of the relata in the intentional relation. In particular, my focus is on whether propositions can serve as the X of the intentional relation thought to hold between *non-derived* intentional states and X. It could turn out that the intentionality of words/sentences, for instance, cannot be accounted for without invoking propositions. But my focus is not on words/sentences.

That being said, the first question to ask is whether propositions are concrete or abstract. Lewis, though sceptical about the clarity of the concrete/abstract distinction, takes possible worlds to be as concrete as the actual world. But as "the objects of thought in general are not sets of possible worlds," (Lewis, 1986, pp.55) taking propositions to be concrete in this sense precludes them from consideration in the present context.[23] Might there be other ways of construing propositions to be concrete? Dummett notes that a subset of propositions—the true ones—are sometimes identified with facts (Dummett, 2006). If we take a fact to be concrete—as, perhaps, a state of affairs (see Wittgenstein, 1922/1961)—then at least some propositions are concrete.[24] The issues here run deep, but this brief sketch of how some propositions (the true ones) can be conceived of as concrete leaves precisely half of all propositions (the false ones) hanging in the balance. On the other hand, there might be a class of propositions that are concrete in virtue of having concreta as constituents. Singular propositions, for instance, might be thought of as inheriting concreteness from that of their concrete constituents. Again, however, since not all propositions are singular, not all propositions are

---

properties on Lewis' view, they are not generally the objects of thought, since the objects of thought are not sets of possible worlds.

[23] Assuming, that is, that Lewis is right.

[24] Although this may conflict with the desideratum that propositions be bearers of truth/falsity. Concrete states of affairs do not seem like the kinds of things capable of truth and falsity.

concrete in this sense. In any event, any view that admits of propositions, will, it seems, have to allow that some are abstract.

Returning to the question at hand: Is intentionality a relation to propositions? If the answer is an unqualified 'yes', then it follows straightforwardly that we cannot have sub-propositional intentional states. This view, call it *propositionalism,* seems problematic (see, for instance, Montague, 2007). For one thing, it seems to rule out anything incapable of propositional thought from having intentionality. Small children, certain mammals, and so on, will all be incapable of aboutness. Second, it seems to deny that we can merely entertain a thought about X without thinking anything in particular about X. But this kind of thing seems to happen all the time. Sometimes, random thoughts about a person/object just pop in and out of my mind, without my thinking anything more about that person or object. Just now, I thought to myself 'Spring', smiled, and then refocused my attention on writing. Perhaps someone will object that what I actually thought to myself was something like 'spring is coming...great', hence the smile. However, regardless of the plausibility of this particular propositional reconstruction, I think the burden of proof is on the propositionalist, if her claim is that all such seemingly simple intentional states can be reconstructed propositionally.[25]

What is required here is some account of propositions such that the intentionality of *some* mental states could be a relation to propositions, whereas the intentionality of other simpler intentional states need not be. This will allow that when I think to myself *that spring is finally here* I am bearing the thinking about relation to the proposition 'that spring is here'; and when I merely think *spring*, I am bearing the thinking about relation to

---

[25] See (Montague, 2007) for extended discussion on the merits/flaws of propositionalism.

something simpler than an entire proposition. An added bonus would be for such an account to show how my mere *spring* thoughts are somehow related to my propositional thought *that spring is here.* And indeed, the two most famous accounts of propositions—the Fregean account and the Russellian account—have sought to do that.

Before looking at each, a caveat: My purpose here is to examine whether propositions—be they Fregean or Russellian—are unproblematic candidates for being one of the relata in the intentional relation. My characterizations of each of their respective views will therefore be partly a function of this purpose. That is, if one were to make a list of all the things the Fregean or Russellian holds about propositions, only some of those things will be relevant to the question 'are propositions what the intentionality of mental states relates us to?'. Others will undoubtedly be relevant to questions in epistemology, the philosophy of language, mathematics, logic etc. My characterization of their respective views will focus on those things I take to be relevant to the former question.

Though Fregean propositions are often contrasted with Russellian propositions, the two have more in common than the standard juxtaposition would seem to imply.[26] Both Frege and Russell took propositions to be timeless, abstract, unchanging, mind-independent entities that served as the objects of the propositional attitudes, and were accessible by multiple thinkers at the same time.[27] Both also took propositions to be complexes, the nature of the simpler parts of which is partly what distinguishes their respective views.

---

[26] See (Makin, 2000, pp. 135 – 178) for an excellent chapter on the similarities and differences between Russell and Frege.
[27] I am not alone in interpreting Frege and Russell as Platonists in this way. See, for instance (Dummett, 1991, pp. 249-262), (Dummett, 2006, pp. 8 - 12), (Soames, 2014, pp. 25 - 33) (Makin, 2000, pp. 139 - 150)

For Frege, propositions ("thoughts", see Frege, 1997c) are complex *sinn*, made up of simpler *sinn*. The proposition 'that spring is warm' is made up of the *sinn* of 'spring', 'warm', and so on, and is itself the *sinn* of the expression 'spring is warm'. In short, the simpler parts of a propositional *sinn* are themselves *sinn* (Frege, 1892/1997). On the other hand, the constituents of Russellian propositions could be concrete objects/people, relations, properties and so on.[28] Notice, however, that both views can accommodate sub-propositional intentionality: On Frege's view, my intentional state as of Spring turns out to be a relation to the *sinn* of Spring, and that very *sinn* figures in the *sinn* of the proposition 'that Spring is coming'. On Russell's view my simple Spring-thought bears a relation to the actual season, and that actual season is a constituent of the proposition 'that spring is coming'.

Minimal though this account of the two classical theories of propositions undoubtedly is, we have enough here to pose the question I am interested in: Could Fregean or Russellian propositions be that to which we are related via intentionality? If propositions are Fregean in the sense that their constituent parts are other *sinn* that—along with the propositions of which they are constituents—occupy what Dummett calls a "third realm" (Dummett, 1991), then the concerns I raised in the context of discussing the p-property theory would seem to arise here as well.

First, it is not entirely clear how items that occupy a distinct ontological realm can have a causal influence on the realm of concreta. Frege's answer is

---

[28] In what follows, I will do my utmost to avoid directly addressing the problem of the unity of the proposition. What it means to be a constituent of a proposition, and how a proposition's constituents are bound together, are difficult and important questions, but far beyond the scope of this thesis. Where the topic is broached, I will not focus on the merits of anyone's account of unity. Instead, I will be examining what a particular account of unity implies about the suitability of propositions to be one of the relata in the intentional relation.

that propositions can have causal effects on the *external* (to the mind) world of concreta in virtue of being *grasped* by competent cognizers, who then give the proposition a causal voice:

"If…I grasp the thought [proposition] we express by the theorem of Pythagoras, the consequence may be that I recognize it to be true, and further that I apply it in making a decision, which brings about the acceleration of masses (Frege, 1918/1997, p. 344)"

This seems fairly intuitive. But how exactly does the proposition affect the cognizer in the first place? How does grasping a proposition "bring about changes in the inner world of the one who grasps it" (Frege, 1918/1997, p. 345)? Certainly the cognitive act of grasping the proposition will bring about inner changes in the same way that the physical act of grasping my rum will bring about inner changes in my arm muscles. But this is presumably not what Frege has in mind: The fact that it is rum that I grasp, rather than vodka, orange juice, or water makes no difference to these internal changes. I take it that what Frege is after here is an account of grasping that makes relevant the fact that it is a grasping of Pythagoras', rather than Fermat's, theorem. To put things another way, I take it that in addition to the properties of the grasping, Frege wants the properties of what is grasped to be relevant to the inner changes. And this seems more problematic since propositions and their constituents exist outside of the spatiotemporal realm.

In a similar vein, how does intentionality[29] relate concrete mental states to entities in the third realm? For Frege, "the thought [proposition]…gets

---

[29] I realize that Frege, and probably Russell, would most likely reject this way of speaking, if not the entire line of questioning. Both took propositions to be the primary bearers of intentionality, and the act of grasping said propositions is what imbues mental states with intentionality. The intentionality of propositions is conceptually prior to that of mental

clothed in the perceptible garb of a sentence, and thereby we are enabled to grasp it" (Frege, 1918/1997, p. 328). I find this clothing metaphor a bit cryptic. Does this make our grasping of a proposition dependent on language? However, just a few pages later, Frege says, "[a]lthough the thought [proposition] does not belong with the contents of the thinker's consciousness, there must be something in the consciousness that is aimed at the thought [proposition]" (Frege, 1918/1997, p. 340). While these two quotes are not obviously inconsistent, I am not sure whether the grasping relation in the former is the same as the aiming relation in the latter. On a plausible interpretation—one on which the two relations are distinct—Frege's claim is that we are first able to grasp a proposition via that sentence that expresses it. Then, once the grasping has occurred, our consciousness gets appropriately aimed at the relevant proposition. There are probably any number of ways this might be developed, but again, my concern here is how (or whether) the intentionality of our mental states is a relation to propositions. The worry is that if propositions occupy this third realm, then the relation appears a bit mysterious. The story that propositions are clothed in sentences and thereby serve to aim one's consciousness, is certainly coherent, but it does not alleviate much of the mystery.

On the other hand, though Russellian propositions also occupy a third realm, they can have concrete things as constituents:

"...[I]n spite of all its snowfields Mont Blanc is a component of what is actually asserted in the proposition 'Mont Blanc is more than 4,000

states, or particular acts of grasping/acquainting (see Soames, 2014, pp. 33). But this, according to Soames, is precisely the source of the so called problem of the unity of the propositon. Roughly: You cannot get unity without the act of unifying—an act that *we* perform (by predicating, or some such). But this is precisely the answer that an account that takes propositonal intentionality (and therefore unity) to be conceptually prior to the cognitive acts of grasping cannot give. *We* (or our cognitive acts) unify propositions (Soames, 2014).

meters high'.   We do not assert the thought, for this is a private psychological matter: we assert the object of the thought, and this is, to my mind, a certain complex (an objective proposition, one might say) in which Mont Blanc is itself a component part" (Russell, 1904/1988, p. 57).

Again, what I am interested in here is whether propositions can be that to which we are related via intentionality. The main concern about Fregean propositions is that it is a bit mysterious how abstract entities composed of other abstract entities, all of which exist in a third realm, can play this role: how does intentionality relate concrete mental states to these kinds of propositions. On the other hand, because Russellian propositions can take concrete things as constituents, the relation might seem initially more plausible: intentionality might be a relation between concrete mental states and abstracta, but these abstracta can have concreta as constituents. However, it is this very feature of Russellian propositions that raises some concerns in the present context.

One question that arises for Russellian propositions is whether they can adequately cope with intentional inexistence. How can a term for something non-existent contribute what it is supposed to—namely, the thing itself—to the proposition in which it occurs?   Russell's view appears to change significantly between 1903 and 1905:[30]

"Whatever may be an object of thought, or may occur in any true or false proposition, or can be counted as *one*, I call a term…I shall use as synonymous with it the words unit, individual, and entity. The first two emphasize the fact that every term is *one*, while the third is derived from the fact that every term has being, *i.e. is* in some sense.   A man, a

[30] *Principles of Mathematics* was originally published in 1903, and *On Denoting* in 1905.

moment, a number, a class, a relation, a chimera, or anything else that can be mentioned, is sure to be a term…" (Russell, 1903/1996, p. 43).

After distinguishing two kinds of terms, things and concepts, Russell continues:

"Points, instants, bits of matter, particular states of mind, and particular existents generally, are things, and so are many terms which do not exist, for example,…the pseudo-existents of a novel" (Russell, 1903/1996, p. 45).

This suggests that Russell endorsed a kind of Meinongianism (see Quine, 1966) at the time (i.e. in 1903), and that his views shifted significantly by 1905.[31] If this is the correct interpretation of (early) Russell, then his answer to the problem of intentional existence would be to deny 1) below.

1) I can think about things that do not exist.
2) I cannot be related to something that does not exist.
3) Thinking about something = being related to it.

Again, we are assuming that Russellian propositions take *terms* (in the sense above) as constituents.   Therefore, if intentionality is a relation to propositions, and propositions have terms as constituents, then intentionality is also a relation to the terms of a proposition. If I think to myself that chimeras are green, then I bear an intentional relation to chimeras. (Early) Russell's (and Meinong's) solution to this is to deny that chimeras do not exist, and therefore deny 1). While this is a possibility we did not consider in the section on intentional inexistence, the problems with

---

[31] However, see Makin (2000, pp. 52 – 57) for arguments to the contrary.

such a view have been well documented (by Russell himself, see Russell, 1905/2008). I will not rehearse these here, as they focus mainly on the ontological/logical problems that follow from countenancing such things, and my concern is what sorts of things intentionality relates us to. Instead, I will simply note that on such a line, intentionality appears to be a rather fantastical relation: in one instance it relates us to quite ordinary things (propositional constituents) such as people, and in the other it relates us to quite extraordinary things like chimeras. At the very least, that seems like a rather unusual kind of relation.

In any event, the theory of descriptions that first emerges in (Russell, 1905/2008), changes tack (Makin, 2000, ch. 3 notwithstanding) on the existence of chimeras, unicorns and the like. Perhaps from a semantics of natural language point of view the theory of descriptions appears promising, from a phenomenological point of view, it appears less so. When I think about Pegasus, my thoughts appear to be about the concrete, flesh and blood (though non-existent) Pegasus. What they do not seem to be about are descriptions. I certainly do not intend this to be an argument against Russell's theory of descriptions. My point is merely that if intentionality is a relation to Russellian propositions, and if the occurrence of 'Pegasus' in such propositions is to be analysed away by the theory of descriptions, then this account of intentionality appears phenomenologically inadequate.

Finally, I want to discuss a feature of Russellian propositions that I find troubling in the present context of trying to find the right relatum in the intentionality relation. However, to bring out the concern I cannot avoid broaching the issue of the unity of the proposition. Soames (2010, pp. 11-

32)[32] notwithstanding, the problem of the unity of the proposition is the problem of saying how the constituents that make up a proposition manage to adhere together, and thereby make it a proposition at all. For Frege, the glue that binds the constituent *sinn* into a cohesive complex *sinn* (thought/proposition) is the saturation, or lack thereof, of the constituent *sinn* (Frege, 1892/1997). It is the unsaturatedness of certain *sinn* that draw the saturatedness of other *sinn* towards them, therefore binding the proposition.[33] For Russell, the relation indicated by the verb binds propositions:

> "Consider, for example, the proposition "A differs from B". The constituents of this proposition, if we analyze it, appear to be only A, difference, B. Yet these constituents, thus placed side by side, do not reconstitute the proposition. The difference which occurs in the proposition actually relates A and B, whereas the difference after analysis is a notion which has no connection with A and B...A proposition, in fact, is essentially a unity, and when analysis has destroyed the unity, no enumeration of constituents will restore the proposition...The verb, when used as a verb embodies the unity of the proposition..." (Russell, 1903/1996, pp. 49-50).

Along with King, (King, 2007, p. 23), I take this passage to indicate that the relation that the verb contributes to the proposition is what does all the work. Remember that on the Russellian view, the constituents of the proposition 'Mount Gugu differs from Mount Kilimanjaro' are the actual

---

[32] See n. 28 above. In fact, I take Soames' point to be that there is a deeper *source* of the problem of unity, not that the problem of unity is not the *real* problem.
[33] Please excuse my rather crude explication of Frege here. I have always thought of the binding of saturated and unsaturated *sinn* to be a process similar to the osmotic activity of roots. When the amount of mineral salts in the roots of plants is less than that of the surrounding soil, the roots draw in the water containing the salts.

mountains and the actual relation of difference that holds between them. It is this relation that unifies Gugu with Kilimanjaro into the proposition 'Mount Gugu differs from Mount Kilimanjaro'. King has his own concerns with Russell's account, the most pressing of which is that this appears to collapse the distinction between propositions and the facts that make them true, which in turn, makes problems for the possibility of false propositions (see King, 2007, p. 23).

My concern is different. Recall that the question I am concerned with is whether Russellian propositions can be that to which we are related via intentionality. This detour into Russell's account of the unity of the proposition has shown that what makes a proposition a proposition—i.e. what unifies its constituents so as to form the proposition—is the relation that the verb contributes. If propositions are the things to which we are related via intentionality, then what makes something a proposition, what unifies its parts into a whole, is extremely important. For Russell, what makes 'Mount Gugu differs from Mount Kilimanjaro' a proposition is the relation that holds between Mount Gugu and Mount Kilimanjaro—the relation that the verb 'differs' contributes to the proposition. And that is what gives me pause. A relation that holds between two things on the other side of the world makes the proposition 'Mount Gugu differs from Mount Kilimanjaro' a possible relatum of my intentional mental states. How can a relation between two things on the other side of the world have this kind of effect on what I am intentionally related to? At least on a view like Frege's, where it is the (unsaturated) *sinn* of 'differs' that creates the proposition that acts as the relata of my intentional state, I need only grasp the proposition's constituent *sinn*. But on Russell's view, the relation indicated by the verb 'differs' unifies the proposition by relating its other constituents, namely Gugu and Kilimanjaro. But how can such a relation accomplish this; how

does a relation that holds between two things on the other side of the world make possible my thinking the proposition 'Mount Gugu differs from Mount Kilimanjaro'? To put things another way, if Gugu and Kilimanjaro do not stand in the difference relation, then there is no proposition, and it therefore cannot be that to which I am intentionally related.

To summarize, the view we have been considering is whether propositions can be that to which we are related via intentionality. After pointing out that what Montague calls propositionalism (Montague, Against Propositionalism, 2007), seems too exclusive for an exhaustive theory of intentionality, we examined both the Fregean and the Russellian views of propositions. As occupants of what Dummett calls the third realm, Fregean propositions faced similar challenges to the p-properties view discussed in section 2.2.1.: How can items that occupy a distinct ontological domain have a causal influence on the realm of concreta? Frege's answer is that they can have such influence in virtue of the causal input of competent cognizers/graspers. But this raised the question of how exactly the propositions affect the cognizers in the first place, and whether the changes in the "inner world" of those who grasp the relevant propositions would be the kind of changes Frege is after (see p. 37 above). We also raised some concerns about how intentionality could make contact with this third realm, and suggested that the idea that propositions get clothed in sentences does not help resolve the mysteriousness of the view. Prima facie, Russellian propositions seem less mysterious candidates for relata of the intentional relation on the count of (sometimes) having concreta as constituents. We raised the problem of intentional inexistence for the constituents of Russellian propositions, and noted that both of Russell's solutions—Meinongianism and the theory of descriptions—were not without difficulties, though for different reasons. Finally, I suggested that Russell's account of the unity of the proposition makes propositions less than

ideal candidates for that to which we are related via intentionality .

To be sure, I have not provided any knock down arguments against the possibility of Fregean or Russellian propositions being that to which we are related via intentionality. What I *have* tried to do is to explain that on the two most prominent views of propositions, propositions are not unproblematic candidates for playing this role.

2.3 Hybrid views

What I have considered so far are cases where the intentional relation relates us to one sort of thing, or another (but not both). I have not considered hybrid views, on which intentionality is sometimes a relation to, for instance, ordinary concrete objects, and other times, abstracta such as properties or propositions. The Impure DR-relation view, mentioned above (section 2.1.4), would count as one such view. Briefly, on the most encompassing hybrid view, my Sibyll-thoughts bear the intentional relation to my flesh-blood Pyrenees, my square-root-of-pi thoughts bear the intentional relation to an abstract number, my belief that spring is right around the corner to a Russellian proposition, and my belief that unicorns have horns to a Fregean proposition. Undoubtedly, this is the most pretheoretically plausible relation view: Pretheoretically, it seems like I can indeed bear the thinking about relation to all sorts of things, such as numbers, properties, propositions, and so on. However, one wonders how to decide when intentionality is relating us to one sort of thing and when it is relating us to another. Obviously, there will be unproblematic cases, such as when I am thinking about my left hand, and when I am thinking about the lowest prime number. But to handle cases of intentional inexistence, representation under an aspect/description, and

46

so on, by arbitrarily stipulating that such cases are all ones where the intentional relation is a relation to some entity of an ontological category that can countenance such things, seems a bit like a cheat.

More importantly, one wonders whether a single relation can play all these roles. Causation, for instance, appears only to hold between entities of the same ontological category. Being taller than also appears incapable of relating something concrete to something abstract. Perhaps there is a class of relations that can traverse the divide between ontological categories, such as relations to numbers. If x is an ordinary concrete object, then it seems inoffensive to say that x bears a relation to an abstract object—namely the number corresponding to its mass. However, the central concern here is not the possibility of there being relations between entities of different ontological categories—though, as the previous few sections point out, this may indeed be problematic—but the possibility that a single relation could hold between a mental state and a p-property in one instance, an a-property in another, a proposition in another, and so on, seems like a lot of work for said relation. Perhaps intentionality *is* a unique relation that can play all these roles. I have not considered such a view, and I think that for someone genuinely intent on providing a hybrid view, this is probably the best option. On the other hand, for someone intent on one of the views we have already considered, I think the most plausible option is a Fregean type view. At least on such a view intentionality is a univocal relation in the sense that it is always a relation to one ontological realm—namely the realm of *sinn.*

At this point, a brief summary is in order. We began by noting that if intentionality is a relation, then two questions emerge: 1) What kinds of things does intentionality relate us to, 2) What kind of relation is intentionality. This chapter has focused on 1). I do not pretend to have

proven or disproven anything. Rather, what I take myself as having accomplished is raising concerns that arise for many of the most popular answers to 1). I have examined some concerns that crop up for views that take intentionality to be a relation to ordinary objects—these stemming mainly from problems about intentional inexistence. I have also mentioned some worries for views that take intentionality to be a relation to properties—be they Platonic or Aristotelian. I considered the possibility that propositions, whether Fregean or Russellian, are one of the relata in the intentional relation, and found some difficulties with both construals. Finally, I considered the possibility of providing a hybrid view, but noted that such a view demands quite a lot of a single relation.

In the end, many of the concerns I raised for the various views I considered revolved around the slightly mysterious nature of the relation that would be required to relate mental states to the favoured entity, or entities. I thus set aside the issue what kinds of things intentionality relates us to, and focus instead on the nature of the intentional relation.

Chapter 3: The Intentional Relation

At this point, I propose to set to one side the issue of what kinds of things we are related to via intentionality, and look instead at several attempts to specify the intentional relation. The goal, once again, is to raise some concerns about the relational view. I mentioned earlier that if we assume the relation view, then two questions emerge: 1) What sorts of things does intentionality relate us to, and 2) how exactly are we to understand the intentional relation? The former was the focus of the preceding chapter. The latter is the focus of the present chapter. The idea is that if there are concerns about the leading answers to both questions, then it might be worthwhile to look into alternatives to the view that prompted both questions—namely the relational view of intentionality.

To the extent that the following various accounts avail themselves of everyday or abstract objects, the worries I raised in the last chapter apply here. In what follows, I examine several leading views of the intentional relation, and raise familiar, sometimes perennial, problems that crop up for each. None of the concerns I raise are intended to be conclusive refutations. Instead, I want simply to rehash some old complaints and raise a new one here and there. My goal is to remind the reader of some concerns regarding the views under consideration, in the hopes of motivating an alternative thesis that I will examine in Part 2.

By far the most popular and widespread relation views are tracking based accounts, e.g. (Dretske, 1981), (Dretske, 1997), (Fodor, 1992), (Millikan, Biosemantics, 2002). Though they differ in the details, all share some basic features. First, and most obviously, all take intentionality to be a kind of tracking relation, where tracking is something like keeping track of or

indicating something.[34] Second, most aim at providing a naturalistic account of intentionality. While different theorists have different views about what naturalism amounts to, some common criteria include the requirement that intentionality be the kind of thing that could be had by a purely physical entity; and that no semantic/intentional terms be invoked in the analysis.[35] Third, with few exceptions, e.g. (Millikan, 2002), the majority of tracking views take causation to figure centrally in tracking.

3.1 Causal Based Tracking Theories:

The initial impetus for causal based tracking theories can be traced back to Dennis Stampe (Stampe, 1977). Stampe considers a case of photographic representation. Why, asks Stampe, is a picture of one identical twin a representation of that particular twin and not her sibling? Resemblance cannot help us here: the twins are identical in appearance. According to Stampe, the reason is that one twin figures in the causal history of the picture, whereas the other does not. Causation, it appears, plays a central role in the representation relation.

We might thus construct a crude theory[36] of intentionality based on Stampe's observation: Mental state type, M, is about some external world item, X, because tokens of X cause tokens of M. Let us assume that M is the mental representation, concept, or mental picture of a dog; and X is a dog, the property of being a dog, or etc. The reason, according to this crude theory, that M is about dogs is that dogs (or the property of being a dog) cause it. However, the above view suffers from a well-known defect, and looking at

---

[34] A compass tracks magnetic north, radar tracks flying objects, and thermometers track the temperature in the room.

[35] This one is especially evident in (Fodor, 1988, ch. 4).

[36] A "crude causal theory" in Fodor's parlance (see Fodor, 1988).

how various theorists have sought to rectify it will help further sort causal-based tracking theories.

The defect is this: There are all sorts of things that can cause me to have thoughts about dogs.[37] Blows to the head, hallucinogenic drugs, coyotes and veterinarians numbers on caller ID can all cause me to have dog-thoughts. Thus, according to this simple theory, my dog-thoughts are not just about dogs, but also about blows to the head, hallucinogenic drugs, coyotes and caller IDs. Moreover, the theory would seem to preclude the possibility of misrepresentation: How can I possibly get things wrong if everything that causes a particular thought in me counts among the items in the thought's extension (Fodor, 1988)?[38] Causation seems insufficient to distinguish the kinds of causes that count as the contents of our dog-thoughts from those that do not. What is needed here is some principle whereby we can sort content fixing causes (hereafter content causes) from other causes—a theoretical supplementation that allows us to rule out bad causes, and say in such cases that one is misrepresenting.

3.2 Normal Conditions:

One such attempt to supplement the causal relation is to specify conditions under which X causing M would suffice for M's being about/meaning X, and holding that these conditions are "normal". Here too, Stampe (1977, sect. 8) seems to be among the first to invoke such conditions within the context of a causal theory.[39,40] If, for instance, we want to specify the normal conditions

---

[37] I use 'thought' here fairly loosely. E.g. I count tokening the mental symbol/concept DOG as a dog thought.

[38] This is a very cursory explanation of what Fodor calls the disjunction problem, but a great deal will be said about it in what follows.

[39] Stampe calls them "fidelity conditions".

under which we represent horses accurately, we would look at the occasions in which horses cause us to have horse-thoughts and specify the conditions that obtain in those occasions as normal. Then, when something other than horses causes us to have horse thoughts, we would find that one or more of the conditions that we have dubbed normal failed to obtain, provided our list of normal conditions was good. Hence, though a cow on a dark night can sometimes cause a horse thought, the fact that the normal conditions failed to obtain allows us to rule out cow from the content of our horse thought, thus delivering a verdict of misrepresentation in such instances.

With respect to intentionality, the normal conditions can be things such as proper lighting, an unobscured view, the absence of an evil genius manipulating your neurons, etc. To use our previous example, your dog thoughts are about dogs because, under normal conditions, dogs cause them. The intentionality of your dog thoughts is explicated in terms of the causal relation that holds between your representing a dog and the presence of dogs under normal conditions.

3.3 Objections and Responses

3.3.1 A Worry About Too Long a List

First, there is a general worry about how to specify all the normal conditions, and what such a list would look like. Recall that it is these normal conditions that are supposed to do the work of weeding out bad causes from content causes so as to allow for misrepresentation. So, with every possible bad-cause, the list of normal conditions would need to include, or be expanded to

---

[40] Color theorists, in the realist camp, often appeal to such conditions as well (see, for instance, Tye, 2002).

include, some condition—deemed normal—that failed to obtain. The idea here is that the list of normal conditions would not just specify good lighting, proper distance, etc. It would also have to include that the representer is not wearing obscuring lenses, is not under the influence of psychedelic drugs, is not seeing a distorted reflection, is not subject to the experiments of an evil genius, is not color-blind, is not pathologically obsessed with seeing cows so as to be suffering from some sort of perceptual self-deception, etc.

3.3.2 The Distinction Between Normal Conditions and Content

Another concern here comes from Fodor (1992, pp. 42 - 43), who asks what exactly keeps the proposed normal conditions from being part of the content, rather than mere determinants thereof. That is, if my dog thoughts are about dogs because under normal conditions dogs cause them, then what makes the content of my dog thoughts dogs rather than dogs-under-normal-conditions?

3.3.3 Do Non-Existents Have Causal Powers?

Fourth, recall that the normal conditions line is still a causal theory. There is thus a question about what to say about things that do not exist. I can have thoughts about thestrals, dragons, Penrose triangles and Golden Mountains. With respect to mere non-existent things, an appeal to counterfactuals might help here: I can represent or misrepresent thestrals because at the nearest possible world where thestrals exist, the normal conditions for the tokening of thestral-thoughts would be thus and so. However, it is not immediately clear how this line would work for impossible objects/properties. Is there some possible world wherein Penrose triangles and round squares are

instantiated such that we could glean the normal conditions for the tokening of the relevant thoughts (see Kriegel, 2011, p. 138)?

3.3.5 An Unnatural Worry

A fifth, more technical objection comes from Fodor (1992, p. 44), who argues that there is a strong teleological factor underpinning normal condition accounts. Writes Fodor:

> Thermometers are OK; given normalcy conditions...the nomic covariance between the length of the column and the temperature of the ambient air determines what the device represents. Violate the normalcy conditions and, intuition reports, you get wild readings; i.e., *mis*representations of the temperature. But, of course, thermometers are *for* measuring something, and precisely what they're for measuring...is what the present analysis treats as a causal (rather than a normalcy) condition. Compare, by way of contrast, the diameter of the coin in my pocket. Fix my body temperature and it covaries with the temperature of the ambient air; fix the temperature of the ambient air, and it covaries with temperature of my body. I see *no* grounds for saying that one of these things is what really represents and the other is a normalcy condition (1992, p. 44).

According to Fodor, what accounts for the difference in the two cases, is that in the case of the thermometer, we have a sense of what it is for. And since "being "for" something is surely a matter of being *intended* for something (Fodor, 1992, p. 43)", the normal conditions line seems to be smuggling in the intentional, pace naturalism. However, if you find yourself within the naturalist-causal-tracking camp, then the thing to do here would be to specify some naturalistically kosher form of teleology.

3.3.6 Intentional Perfection

All that being said, the most pressing concern for the normal conditions line is that it deems it impossible to misrepresent when conditions are normal. That is, if all the normal conditions obtain, then we cannot misrepresent. With respect to this point, I direct the reader to Mendelovici (2010, pp. 28 – 38, MS). I will not rehearse her entire argument here, but I will note that there are several plausible examples (e.g. color and pain representation) where we consistently, and reliably, misrepresent things. And while some of these cases might be ones where normal conditions fail to obtain, it does not seem likely they all are. Otherwise, in virtue of what should we say that the conditions that always fail to obtain in these problematic cases are normal? That is, if the normal conditions that are required for accurate color/pain representations never obtain, then in what sense are they normal? See also (Mendelovici, 2013), who argues that the mere possibility of cases of reliable misrepresentation is problematic for tracking theories such as the normal conditions view.

3.4 Teleology: Evolution and Education

A second[41] way to supplement the causal relation is by appeal to teleology. Again, what I have to say here pertains to causal-teleological theories (e.g. Dretske, , 1981, 1997). Millikan's (Millikan, Biosemantics, 2002) non-causal teleological view will be addressed later. While causal teleological theories are varied, at the core of most is some common ground. Recall our initial

---

[41] I say second here, but a normal conditions advocate might specify the normal teleologically: The normal conditions are those that instantiate a design conditions type— where design conditions are those under which a particular causal-intentional relation was selected for by natural selection.

55

crude causal theory: intentional states track (viz. are about) their causes. This turned out to be problematic because causation alone is insufficient to distinguish between content causes, and other causes. Whereas the normal conditions advocate contended that my dog-thoughts are about dogs (rather than convincing dog statues) because, under normal conditions, they are caused by dogs, the teleological theorist contends that my dog-thoughts are about dogs because it is their job to track, indicate, or provide information about dogs. To be as straightforward as possible, the question is: Given that intentionality is a causal relation, how do we sort out good causes from bad ones? The teleological theorist says that the good causes are those that occur when the thing doing the representing is functioning properly. The causal relation is supplemented by an appeal to proper functioning. The challenge now is to say in virtue of what a particular representation has the particular job/function it does: Why is it my dog thoughts' job to be about, track, indicate or carry information about dogs?

3.4.1 Natural Selection

One prominent answer here is that natural selection determines what job a particular representation has. For example, a rabbit's dog-thoughts have the job of tracking dogs because it was this function that helped its ancestors survive and reproduce. That is, nature favored those rabbits whose dog thoughts preformed this function.

3.4.2 Which Cause?

A set of related concerns here stem from the fact that natural selection does not appear to deliver the kind of content determinacy we might want for our theory of intentionality. For instance, it seems perfectly plausible that

rabbits whose dog thoughts were caused by not only dogs, but also wolves, wolverines, etc., would have been equally favoured by the forces of evolution. In which case, the function of a rabbit's dog thoughts (as determined by natural selection) would be to track not only dogs, but also all these other dog-looking creatures.

Relatedly, for any causal relation between a dog and a dog thought, there is a host of intermediaries in the causal chain. In the case of vision, light reflected off the dog must stimulate the rabbit's retina, and so on. In olfaction, scent particles must travel from the dog through the air and be received and processed by the rabbit's olfactory system. The problem is that any selection for dog caused dog-thoughts is equally a selection for these causal intermediaries. So natural selection appears inadequate to determine that the function of a dog thought is to track dogs rather than reflected light, scent particles, etc. Moreover, natural selection would seem to favour the set of undetached dog parts→dog thought connection every bit as much as the dog→dog thought connection, so there might be an additional worry about determinacy here.[42]

3.4.3 Naturally Selecting Error

A similar concern involves the possibility of naturally selecting for error. Imagine that a small subset of the early humanoids were visually constituted such that upon encountering berries that others would see as red and ripe, they see the berries a sickly yellow color—quite unappetizing. As hunter-gatherers, berries figure importantly in the diet of early humanoids. Now imagine that due to some rare climatic conditions, one particular species of

---

[42]For extended arguments that only a particular kind of intentionality (or a particular point of view on a given content) can deliver determinacy, in the above sense, see (Searle, 1987), (Horgan & Graham, 2012).

plant, whose red berries are highly toxic, has proliferated and outcompeted its evolutionary rivals. In this case, it seems as though evolution would favour the misrepresenters.

3.5 Learning

Another influential answer to the question of what makes something the function of a representation is Dretske's learning model (Dretske, 1981).

> In teaching someone the concept *red*, we show the pupil variously colored objects at reasonably close range and under normal illumination. That is, we exhibit the colored objects under conditions in which information about their color is transmitted…This is why we cannot teach someone the colors if we put the objects 400 yards away…This is why we do not carry out such training in the dark, or under abnormal illumination…In the learning situation special care is taken to see that the incoming signals have an intensity, a strength, sufficient unto delivering the required piece of information *to* the learning subject. If the lights are too dim, they are turned up. If the objects…are too far away, they are brought closer. If the subject needs his glasses, they are provided (1981, pp. 194-195).

According to this model, an M thought has the function of indicating Ms because of a learning period wherein the representer is trained to token M-thoughts in M instances. The crucial part here is that during the learning period, a teacher ensures that the causal correlation of M thoughts to Ms in the fledgling representer becomes increasingly robust by ensuring optimal conditions obtain during the learning period.

3.5.1 Normal Conditions Again?

One worry about this account is its seeming reliance on normal conditions. It is the teacher's job to ensure that e.g. the lighting is good and the relevant objects are not too far way, etc. If these teacher-insured normal conditions do any theoretical lifting, worries similar to the ones raised above (about normal conditions) apply. One worry was that the list of normal conditions would turn out to be rather long, and this would mean a lot of work on the part of the teacher—ensuring all conditions are met during the learning process. Fodor's concerns (1992, pp. 42 - 44) about finding some principled means of excluding the set of normal conditions from the content of a given representation (sect. 3.3.3 above), and about finding some non-arbitrary means of saying what counts as the normal condition versus the content (sect. 3.3.5 above) apply here too. Moreover, the learning model would need to say something about how we come to have concepts with no worldly extension. For instance, how exactly do we determine the normal conditions for the correct learning of the concept dragon?

3.5.2 A Worry About Naturalism

Relatedly, one might ask how the teacher goes about determining what conditions are optimal for learning a particular concept. That is, why would the teacher think that a particular lighting condition would be most conducive to producing content causes of red thoughts? Presumably, the teacher surveys the conditions under which s/he has red caused red thoughts and duplicates them for the student. The problem here is that this move appears to put things backwards. Rather than getting representations from normal conditions, we are constructing normal conditions from representations. What was sought here was a naturalistic theory of

intentionality/representation—a theory that explained the target phenomenon without appeal to semantic/intentional terms. The hope was that by invoking the notion of the function of a representation, and having this function be determined by a learning period in which a teacher ensures the conditions are conducive to learning a particular concept, we could sort out genuine content causes from misrepresentations. The problem, however, is that the conducive conditions appealed to in our explanation are derived from the representations of the teacher, and so we have not done away with semantic/intentional terms after all.

3.5.3 A Worry About Intentional Smuggling

In a similar vein, Fodor too charges this account with smuggling in intentional items (Fodor, 1992, pp. 41-42). Fodor's objection runs thus: Assume the learning period is over, and the new graduate tokens a dog thought as a result of encountering a fox. Given that the learning period has established a law-like connection between dogs and dog thoughts, and this instance is one that fails to instantiate the law, we have a case of misrepresentation: the new graduate has mistaken a fox for a dog. So far, so good. However, given that a fox caused a dog thought at time T (where T=the moment after graduation), it seems likely that it would have caused the same thought at time T_{-1} (i.e. right before graduation). But then what licenses our classifying this event as a misrepresentation? That is, if a fox would have caused a dog thought during the learning period, why is the content of a dog thought *dog* rather than (*dog or fox*)?

One response here, on behalf of the causal teleological theorist, is that had this event occurred during the learning period, the teacher would have corrected the student. But such a response is not open to the naturalist: As

60

mentioned above, Fodor's point about invoking the intentions of the teacher as an essential part of the explanation of how representations get their content appears to apply.

3.5.4 Evolution as an Inappropriate Tool

While highly implausible, someone determined on the learning model might argue that rather than relying on his/her own representations in order to glean the conditions most conducive to learning, the teacher actually gets these conditions by appeal to something like design conditions.[43] On this line, the teacher tries to replicate the conditions under which a particular causal-intentional relation was selected for by natural selection.

Or, as in the case with Dretske's later work (see Dretske, 1997, p. ch.1), we might dispense with the notion of a teacher, allowing that there is a perfectly clear sense in which being *designed for* (i.e. designed to indicate, represent) does not imply a *designer.* In other words, perhaps evolution alone can do the job of determining the function of X. Dretske seems to think so, at least with respect to what he calls "natural representation" (1997, pp. 7-8). For Dretske, the senses, for instance "...have information-providing functions, biological functions, they derive from their evolutionary history (1997, p. 7)." Evolution has imbued the senses with the function of carrying information to the organism whose senses they are. Olfaction, for instance, has the function of carrying chemical information about the environment in which an organism finds itself.

The worry here is an extension of those I raised in the context of natural selection theories. Because natural selection can confer a survival advantage

---

[43] This is not, of course, Dretske's view. I am merely trying to tick off the possibilities.

on misrepresenters, it can likewise determine that the design conditions are those under which the misrepresentation takes place. To use our previous example, the design conditions under which our ancestors were caused to have their cougar thoughts may be precisely those conditions under which they were caused to have cougar thoughts by bears on dark nights. So, design conditions may be ill equipped to establish the law-like correlation between cougar thoughts and cougars required. Again, there are probably a whole host of possible responses here, and I cannot hope to address them all.

However, abstracting a bit, we might say that evolution, as a system designer, does not carve out the evolutionary history of an organism along semantic lines such as satisfaction, accuracy, or truth; it is blind to failures such as misrepresentation and falsity. Evolution is a process concerned solely with continuation and adaptation. If getting things representationally wrong ensures the continuation of a species, then evolution rewards falsity. I am not saying that we could not construct a theory that takes evolutionary success as the mark of veracity, but such a theory would imply a pseudo-Machiavellian semantics—taking evolutionary success as the only justification for our attributions of truth, falsity, accuracy, etc. And I doubt that any of the theories under discussion would endorse such a move. The point, to repeat, is that evolution seems like a tool better equipped to reward adaptability than veracity.

3.6 Asymmetric Dependence

A third and widely discussed proposal for supplementing the causal theory is asymmetric dependence (Fodor, 1992, 1988). Importantly, asymmetric dependence represents a departure from the accounts we have so far considered. While the previous theories tended to look for what goes wrong

in cases of misrepresentation—e.g. the conditions weren't right for good causings, or there was a malfunctioning—asymmetric dependence focuses on the relationship between misrepresentation and representation. The core idea behind the theory is that getting things wrong somehow depends on it being possible to get things right, but not the other way around. Non-content causes thus depend on content causes in a way that content causes do not depend on non-content causes. What does the work of weeding out non-content causes is their asymmetric dependence on content causes. And what makes a content cause the content of a representation is (among other things) its being that upon which non-content causes are asymmetrically dependent.

3.6.1 The Theory

Setting this last point aside, I want to look a little more closely at Fodor's account. According to Fodor my X-thoughts are about Xs if:

1) 'Xs cause X thoughts' is a law.

2) Some X thoughts are actually caused by Xs

3) For any Y such that Y $\neq$ X, if Ys sometimes cause X thoughts, then Ys causing X thoughts is asymmetrically dependent on Xs causing X thoughts (Fodor, 1988, p. 109).[44]

A brief example will perhaps be helpful here. The reason why horse thoughts are about horses is because horses cause them. And any non-horse, (say) a cow on a dark night, that causes horse thoughts does so in virtue of horses causing horse thoughts. If you break the connection between horses and

---

[44] For reasons of terminological consistency, I have continued to speak of X thoughts, rather than tokens of 'X', concepts, or mental symbols. Also, there are two points Fodor bids us to keep in mind. First, these conditions are meant to apply synchronically (1988 pg. 109). Second, the theory is to be understood in terms of the nomic relations among properties (1992 pg. 102). Nothing I will have to say rides on confusing either of these points.

horse thoughts, the connection between cows on dark nights and horses thoughts would also be broken. But break this latter connection and the horse-horse thought connection remains; hence the asymmetry.

3.6.2 A Worry About Scope

First, and foremost, it should be noted that these are meant only as sufficient conditions for aboutness. Unicorns and Penrose Triangles would not, for instance, be objections here on account of failing to satisfy 1) and 2). That said, condition 1) seems to make the theory problematically limited. Since Fodor is adamant that the theory be understood in terms of the nomic relations among properties (1992 pg. 102), and since laws range over properties, not individuals, there will be a lot more than theoretical, logical and vacuous concepts that the theory cannot account for: Thoughts that appear to involve an individual such as singular thoughts about my father, or proper name thoughts, will also fail to fall under the purview of asymmetric dependence. The concern here is thus that the theory might be too narrow in scope to be the full story of my intentional mental life.

3.6.4 A Concern About Exclusivity

There appear to be some difficulties that arise for Conditions 1) & 2). Conditions 1) requires that X's cause X thoughts be a law, and condition 2) requires that at least some X-thoughts be actually caused by Xs. Given this, the theory cannot allow for situations where an X-thought is always caused by something other than X. As I mentioned in section 2.1.1, Mendelovici (2010) has argued convincingly that color-thoughts are precisely the kind of thoughts deemed impossible by condition 2): Our color representations always represent something other than their causes. Color thoughts are most

plausibly caused by light reflectance profiles, but these are not what I represent when I represent color.[45]

Of course, someone might argue that in light of such arguments we should be color realists; but this seems backwards. We would not, for instance, argue that because our theory of intentionality requires the existence of unicorns that we ought to be realists about unicorns. That said, comparing colors to unicorns does seem slightly unfair. In fact, however, the problem for condition 2) is not that color realism is false, it is that condition 2) requires that it cannot be. In effect, condition 2) rules out the possibility of our consistently misrepresenting something (Mendelovici, 2010). Though it would lose some of its literary effect, the point about unicorn realism could be turned around: a theory of intentionality should not require us to deny the existence of unicorns either (See also Mendelovici, 2013, section 6.2).

For clarity's sake: The worry I am trying to raise here is not that asymmetric dependence cannot countenance unicorn thoughts. It patently can. The story about unicorns is that the property of being a unicorn would cause unicorn thoughts in the nearest world where unicorns exist.[46] "There can, of course, be a nomic connection between properties one or more of which is *de* facto uninstantiated...[U]nicorns...*would* be nomically sufficient for 'unicorn'-tokenings if there were any (Fodor, 1988, pp. 163-164 n.5)". In other words, there is a law-like connection between the property of being a unicorn and the thought it would cause were there any unicorns. No, the concern raised by Mendelovici, that I am here echoing, is that as stated, the theory rules out the possibility of there being nomic connection between Ys and X-

---

[45] For arguments that color realism is false, I direct the reader to C.L. Hardin (Hardin, 1988). For arguments to the contrary see (Byrne & Hilbert, 2003).

[46] In the addenda to *Naming and Necessity*, Kripke offers interesting insights about the possibility of unicorns. However, as endorsing his argument requires the acceptance of essential properties, I chose not to rely on it.

representations that is not dependent on any other nomic connection between Xs and X-representations—where Y are not Xs, that is, where X-thoughts always misrepresent Ys.

Again, perhaps this is not so devastating; perhaps, as it turns out, there just aren't any cases of what can be called *nomic misrepresentation.*[47] However, blunted though the objection might be, there is still something concerning about excluding the possibility of such cases *apriori*. What exactly is concerning about such an exclusion? It seems like the existence, or lack thereof, of cases of this kind of misrepresentation is an empirical matter, to be sorted out by investigation not stipulation (see Mendelovici, 2013).

Condition 3) is the asymmetric dependence clause: For any Y such that Y $\neq$ X, if Ys sometimes cause X-thoughts, then Ys causing X-thoughts is asymmetrically dependent on Xs causing X-thoughts. To use one of our previous examples: Though horses sometimes cause me to have cow thoughts, my cow thoughts are about cows and not horses because horses causing cow-thoughts is asymmetrically dependent on cows causing cow-thoughts. Break the cow cow-thought connection, and you thereby break the horse cow-thought connection; but break the horse cow-thought connection, and the cow cow-thought connection remains. It is this condition that is supposed to do the work of weeding out bad causes from content causes.

3.6.5 Pathological Misrepresentation

In (Adams & Aizawa, 1992), a host of pathological cases are submitted as putative counter examples to Asymmetric Dependence. The following can be

---

[47] Nomic misrepresentation = misrepresentation that occurs in a law-like manner.

seen as raising the same sort of concern: Assume that 'dogs cause dog thoughts' is a law, and that some dog-thoughts are in fact caused by dogs, so that conditions 1) and 2) are satisfied. Now, imagine that we find some artificial means of producing dog thought tokens, say a particularly refined blow to the head. The question is whether this artificial means of producing dog-thoughts asymmetrically depends on a law like connection between dogs and dog thoughts. We might even push the example further and say that we have become such accomplished head hitters that we are able to produce, with law-like consistency, dog-thoughts by this means. The question is whether the blow to the head→dog-thought connection depends, in any way, on the dog→dog-thought connection. In other words, and at the very least, it does not appear that the reason why the blow to the head causes dog-thoughts is that dogs do.

Someone keen on asymmetric dependence might argue that extraordinary though it seems, breaking the dog dog-thought connection would in fact break the blow-to-the-head dog-thought connection: If dogs do not cause dog-thoughts, then in virtue of what should we call such thoughts dog-thoughts. The idea, I take it, is that it only makes sense to call a painting a dog painting because of its connection (in this case similarity) to dogs. If this connection is broken, say, because dogs look a lot like ants, then in virtue of what should we call this painting of a large furry animal a dog painting?

The problem with this kind of response is that it seems circular, assuming that concepts are individuated by their referents, which on Fodor's view are their contents. Asymmetric dependence is supposed to give us an account of the aboutness of our X-thoughts—viz. why this particular thought is about Xs. In the process of explaining the relation in virtue of which X-thoughts are about Xs, it seems circular to assume that the thoughts in question are the

thoughts they are in virtue of the relation you are attempting to explain. If my dog thought is the thought it is in virtue of its connection to dogs, then it follows trivially that it is no longer a dog thought if this connection is broken. And if the relevant thought is no longer a dog thought, then it follows trivially that anything else connected to it (say, a blow to the head) will also fail to be connected to a dog thought.

Of course, Fodor does not argue this way. For Fodor, concepts and their ilk are not wholly individuated by their contents, but partially by their form (2008, p. 75).[48] This is how a purely referentialist semantics avoids Paderewski cases:[49] It is because John has two concepts—Paderewski$_1$ and Paderewski$_2$—that he can wonder (coherently) whether Paderewski is Paderewski. And, this being the case, our blow-to-the-head problem remains: Assume that dog-thoughts have the form F, and that the power of dogs to cause dog-thoughts is reducible to their power to produce thought tokens with the form F. A blow to the head's power to produce thoughts with the form F does not appear to asymmetrically depend on the power of dogs to cause such thoughts.

Again, someone keen on asymmetric dependence will undoubtedly be able to add to the theory to blunt some or all of the above objections. That is, I do not think I have provided any knockdown arguments here. What I have tried to do was show that, as a means of supplementing the causal relation in order to allow for misrepresentation, asymmetric dependence is not without some challenges.

3.3.6 Teleology without Causation

---

[48] A similar line is taken by Fodor in his reply to Block (see Fodor 1992 pp.111-112).
[49] Cases, that is, where the same man, Paderewski, is known by some for being an accomplished musician, and by others as a politician.

In her Biosemantics (Millikan, 2002), Ruth Millikan proposes a non-causal based teleological theory. For Millikan, intentionality is still a kind of tracking relation that holds between intentional states and their contents. What is different on her account is that contents are not causally determined. An intentional state is about whatever item/state of affairs in a representation consumer's environment is required for said consumer to function properly—where functioning properly is evolutionarily determined. It is how a consumer uses the representation to function properly that determines the content; and it is how a consumer's ancestors used representations of that type to survive and reproduce that determines what it is to function properly. A consumer, on this view, is a system that exploits the representation in the performance of its proper function. In what follows I will focus on organisms as consumers, rather than distinct subsystems of said organisms. The distinction will not make much difference to what I have to say.

The well-worn case of the frog snapping at flies will be useful here. To find out the content of the frog's representational state when it snaps at a fly, we first ask how ancestral frogs would have used such a representation to survive and reproduce. Presumably, the representation would have been used for the acquisition of energy/nutrition. So the proper function of the frog's representational state is to track sources of energy/nutrition. When a modern day frog snaps at a fly, the content of its representational state is 'source of energy/nutrition'. Notice how well this account avoids the kinds of bad-cause problems associated with causal theories. A causal theory has to avail itself of the kinds of causes that can, in this case, figure in visual representations—e.g. things like shape and color. And it is because of this that a causal theory has a hard time ruling out BB's from the content of the frog's fly representations. On the other hand, Millikan can grant that all sorts

of things can cause a frog to token a fly-thought, because on her view, causes are irrelevant to content. For Millikan, what matters for the determination of content is how the representation is/was used to contribute to the survival and reproduction of frogs.

3.7.1 Why Jack and Jill went up the hill

A set of concerns here stems from the implausible intentional explanations this account seems to deliver in certain cases. In his famous tale of the kimu and the snorf, Pietroski too stresses the implausible nature of the intentional explanations the consumer-based account appears to deliver (Pietroski, 1992). The kimu are a docile species of herbivore that live near a large hill. Their only predators are the carnivorous snorfs, who roam past the hill each morning. At some point in the evolutionary history of the kimu, a particular kimu, Jack, underwent a genetic mutation that caused him to token a particularly pleasant mental state, M, in the presence of the rising sun. Each morning when the sun came up over the hill, Jack tokened M, and ascended the slopes in pursuit of the pleasant light of the early morning sun. As time passed, those of Jack's descendants who inherited the gene flourished as a result of being at the top of the hill, pursuing the pleasant red light, when the hungry snorfs past each morning to consume their non-M-tokening kin.

The rub is that according to Millikan's account, the content of the kimus M-thoughts turns out to be something like 'snorf-free zone', or 'safety this way'. All the ingredients are here: We have some representation consumers (Jack's descendants), who have some intentional mechanism whose production of R-thoughts co-varies with the presence of some environmental feature. Some ancestral kimus were able to survive and reproduce in virtue of using the thought tokens of the M-type to avoid being preyed upon. This determines

70

that the proper function of this intentional mechanism is to direct kimus out of harm's way. Thus the content of the kimus' R-thoughts is 'snorf free zone'.

However, according to Pietroski, though there are several plausible candidates for the contents of the kimus' R-thoughts—e.g. 'sun over there', 'pleasant light that way', etc.—R-thoughts are certainly not about snorfs. Millikan's account delivers the wrong intentional explanation; namely that the kimus are thinking about snorfs when they token M. We are invited to test this by imagining how the kimu would react if a herd of snorfs who had undergone a genetic mutation so as to cause M-tokenings passed by the hill (Pietroski, 1992, p. 276). The intuitive answer here is that the kimus would head on over to the M source. Worse still, it is not so far fetched to think that after several generations of climbing the hill before the immanent arrival of the snorfs, the red loving kimus would have no idea what snorfs are, having never encountered them. And yet Millikan's theory still predicts that the content of the kimus R-thoughts is 'snorf free zone'.

The big picture here is that Millikan's account does seem to capture an important relation between an organism's mental states and items in the environment. After all, the co-variation of M-thoughts with snorf-free zones is certainly useful in explaining the evolutionary success of the kimus. The problem is that this relation looks like a poor candidate with which to identify intentionality. The chief task of any relation view is to specify a relation that explains why a mental state has the content it does. The problem with Millikan's account is that the relation it homes in on does not explain why a given mental state has the content it does, but rather, why a given mental state was of an evolutionary advantage. And though the two explanations can appear to converge—hence why consumer based theories look plausible in some cases—their target phenomena are distinct.

The kimu survived and reproduced because their R-thoughts led them out of harm's way. The proper function of the consumers of M-thoughts, in this evolutionary sense, is to keep kimus from being eaten. This proper function relation is the one Millikan's theory identifies. However, the present case shows that this cannot be the intentional relation, since the content of the kimus' M-thoughts is not 'snorf-free zone'. Kimus could have M-thoughts, and thereby pursue M-ly things, in a world without snorfs.

3.8 Concluding Remarks

We began with two questions that emerged as a result of assuming that intentionality is a relation: 1) what kinds of things does intentionality relate us to? And 2) what kind of relation is intentionality? I considered some of the most popular answers to 1): Ordinary concrete objects, abstract properties and propositions. Ultimately, I concluded that there are concerns about each, and that taking intentionality to be a relation to one, several, or all of these things can make the intentional relation appear a bit mysterious. As said, I do not take myself as having conclusively ruled out any particular view, but merely as having raised some concerns about the most popular views.

Setting the question of what objects might serve as contents aside, I went on to look at a prominent family of views about the nature of the intentional relation. Each attempted to account for intentionality in terms of a list of relational ingredients. These included evolution, learning, asymmetric dependence, and biological proper functioning. The hope of the various theorists considered was that the correct assembly of some of these relational ingredients could transform the favoured relation into intentionality.

The majority of the views considered were causal theories, which I taxonomized by how each attempted to handle the problem of error. Causation is, as it were, blind to semantic-like notions such as truth, falsity, accuracy and error. Intentionality, on the other hand, is a phenomenon for which such semantic notions appear to matter. So the challenge for the causal theorist was to specify some supplemental element, X, that when combined with causation, delivered a semantic relation. My conclusion was that no theory was without its difficulties.

The last view I considered was Millikan's teleological, consumer-based, account. In general, theories that appeal to evolution in their attempts to spell out the intentional relation may seem to have an advantage over rival theories. To repeat, the common goal of all the theories we have considered is to specify how the right combination of non-intentional elements could give rise to intentionality; and the problem is that it is difficult to see how we could get a semantic phenomenon out of a non-semantic relation such as causation. The advantage of evolutionary theories is that the concept of evolution comes, as it were, pre-furnished with normative notions such as success and failure. And because normativity is also a property of those troublesome semantic notions such as truth and falsity, perhaps evolution is precisely the ingredient needed to get intentionality from non-intentional elements. Perhaps we could cash out truth and falsity, or accuracy and error, in terms of evolutionary success and/or failure.

The concern I raised is that the mapping between the concepts of evolutionary success/failure on the one hand, and accuracy and error on the other is not without its difficulties. The Kimus enjoyed evolutionary success because their R-thoughts led them to snorf-free areas. But this evolutionary

success does not translate unproblematically into truth or accuracy. If the snorfs somehow learned to climb the hill, then the very same R-thoughts would spell evolutionary failure for the kimus. But in both cases, the kimus' thoughts were plausibly taken to be about red things.

To reiterate, I do not pretend to have refuted any particular theory about the objects of intentionality, or the intentional relation. What I have done is raised some concerns about many of the most popular answers to questions 1) and 2) above. This prompts the question: Is there an alternative to the assumption that gave rise to 1) and 2)—namely, that intentionality is a relation? It is to this that I now turn.

Part 2 Intentionality and Phenomenality

4. Introduction: Phenomenality and Intentionality

I ended part 1 on an anticipatory note, suggesting that it is worthwhile looking at alternatives to what I called the relational view of intentionality. In this part of the thesis, I plan to do exactly that.

Employing Kriegel's terminology of intentionality being "injected into the world" (Kriegel, 2011, p. 17), we might take the view we are setting aside—the relational view—as claiming that intentionality is injected into the world by a relation to one or more types of things, a causal relation in most cases. In this section, I will explore the alternative view that intentionality is injected into the world in virtue of its close connection to phenomenal consciousness.[50]

As I will use it, 'consciousness' denotes the same phenomenon as expressions such as 'what-it-is-like', 'phenomenality', 'how things are for me-ness', 'how it is for me-ness', 'what it is like for me-ness', 'phenomenal consciousness', 'phenomenal character', 'experience' and 'qualitative character'. With such a myriad of terms, it would be useful to pick one and stick with it; and, as far as I can do that, I will. My preference is for 'phenomenality', but there will be times when it will serve practical purposes to employ the slightly more clumsy expressions 'what-it-is-like' and 'how things are for me' (and their cognates). In any case, all such instances should be taken as denoting one and the same phenomenon, namely phenomenality.

---

[50] The vagueness of idioms such as "close connection to" is deliberate. The view I will examine will become more definite in what follows.

75

Why is it that I think the intentional and the phenomenal are closely connected? First, they are both aspects of mentality. Both often happen together, and, perhaps slightly more contentiously, both appear to converge in several places: perception, olfaction, and so on. Representationalists (intentionalists) such as (Tye, 2002), (Lycan W. G., 1996) and (Dretske, 1997) have noted as much, arguing that at least some forms of phenomenality, such as perception, are intentional. Additionally, it has been noted by several important historical thinkers that consciousness (phenomenality) is always consciousness of, where the 'of' is that of intentionality (Husserl, 1913/1998) (Brentano, 1887/1995).[51]

For Husserl, phenomenological investigation involves bracketing off certain elements (a process called *epoche*) of an experience in order to describe the experience only in terms of what is available from the first person perspective. And while the phenomenological question to be addressed will partly determine what exactly gets bracketed off (Husserl, 1913/1998), the positing of the existence of the things to which an experience is purportedly directed always figures in (Husserl, 1900/2001). That is, in phenomenological investigation, we are to refrain from "naively positing the existence of the objects," or otherwise "going on to characterize them" (Husserl, 1900/2001, p. 170).[52] This follows from Husserl's requirement that the phenomenologist is to describe things exactly how they are from the first person perspective: According to Husserl, from the first person perspective, the existence, or lack thereof, of the purported objects of experience cannot be determined. [53] Assuming one has bracketed off the relevant

---

[51] See (Searle, 1992) for a dissenting view.
[52] Excuse the rough-and-ready characterization of Husserllian phenomenology.
[53] Admittedly, this is a very rough characterization of Husserl's phenomenological approach. For a more thorough discussion of Husserl and

presuppositions, one arrives at an irreducible "abstract structure by virtue of which the mind is directed" (Dreyfus & Hall, 1984, p. 2).

While my understanding of Husserl is rudimentary at best, the point here is not to deliver an exacting exegesis of Husserl, but rather to draw on what I take to be some of his insights. First, both Husserl and Brentano assumed that phenomenality and intentionality were, at the very least, inseparably intertwined. For Husserl, directedness just is a phenomenal feature (Gurwitsch, 1984) (Follesdal, 1984). Second, regardless of whatever else he thought we could glean from phenomenal reduction, Husserl thought that this phenomenal directedness was certainly accessible in this way. That is not to say that Husserl thought phenomenal directedness was obvious and apparent to anyone who gave it a minute's reflection. That it required phenomenological reduction to bring out indicates that Husserl may have thought that the phenomenal directedness was introspectively unobtrusive, subtle. Finally, Husserl's view locates the source of directedness (intentionality) in phenomenality. For Husserl, it is consciousness that injects intentionality into the world.[54]

Earlier, I listed several reasons for thinking that intentionality and phenomenality are closely connected. Having mentioned what I take to be some important Husserllian insights, I want to suggest a final motivation I have for thinking that there is a close connection between intentionality and phenomenality. Like Husserl, I think that careful reflection on some intentional experience can reveal a phenomenal feature of the experience

his contributions to contemporary philosophy of mind and cognitive science see (Dreyfus & Hall, 1984).

[54] Again, I am not hereby endorsing all of Husserl's views. I would not, for instance, follow Husserl in taking the noema—the entities in virtue of which we are phenomenally directed—to be abstract entities (see Follesdal, 1984).

that is best described as directedness, ofness or aboutness. To avoid inviting talk of *what* this phenomenal feature directs us to, or is of/about, we can describe it as a phenomenal purporting-to-be-directed (of/about). Undoubtedly, this feature will be subtler in certain cases, yet no less present. Again, this is not intended as some form of argument, or as evidence for some further conclusion. It is merely an account of one of my principal motivations for exploring views that accept a strong connection between intentionality and phenomenality.

The proposal that intentionality and phenomenality are closely connected stands in need of clarification. First, as a reminder, I am interested in whether *original[55] intentionality* is closely connected to phenomenality. More importantly, what exactly is the nature of this close connection? Undoubtedly there are any number of ways the connection might be construed. It might turn out that phenomenality is identical to, reducible to, or in some way dependent on intentionality. This kind of view is typically associated with representationalists such as Tye (1995), Lycan (1996) and Dretske (1997). However, most representationalists take intentionality to be the kind of causal/covariational relation we discussed in chapter 3 (see, for instance, Tye 2002; Dretske, 1997). The thinking is that intentionality is reducible to some naturalistically acceptable causal/covariational tracking relation; phenomenality is reducible to intentionality; therefore phenomenality is reducible to some form of naturalistically acceptable causal/covariational tracking relation. As one of our central purposes here is to investigate non-relational alternatives to the relation view of

---

[55] Hereafter, all uses of 'intentionality' (and its cognates) are to be understood as elliptical for 'original intentionality'.

intentionality, the representationalism espoused by Tye, Lycan and Dretske, do not fit the bill.[56]

On the other hand, it might turn out that intentionality stands in a dependence relation to phenomenality—perhaps supervening on it. Intentional content would then be determined (at least partly) by phenomenal character. Call this the *supervenience view*. Or, someone might take a very strong line on the relation between intentionality and phenomenality—one on which intentionality and phenomenality are type and token identical. On such a view intentional content and phenomenal character name the same phenomenon. Call this the *strong identity view*. Finally it might be that intentionality is a kind of phenomenality. On such a proposal, intentional content would be a kind of phenomenal content. Call this the *moderate identity view*.

In the next chapter, I will explore possible ways a theory of intentionality might be developed in accordance with what I take to be the most plausible of the three abovementioned views—the moderate identity view. I set aside the supervenience view for the following reason: It is not obvious to me that a relational phenomenon could not supervene on a non-relational phenomenon. The grasping relation I bear to my coffee appears to supervene on the neuromuscular events in my arm. If this is the case then, the view that intentionality supervenes on phenomenality may not be a non-relational view after all. I also set aside the strong identity view. This is due mainly to the fact that the strong view has certain implications that stand in need of more time and space than I can reasonably allow for, given the time and

---

[56] That is not to say that representationalism per se is inconsistent with a non-relational view of intentionality: One might argue that phenomenality is reducible to intentionality, and that intentionality is not a relation. I set this proposal aside for the time being.

space I devote to what I will later call the problem of cognitive contact. For clarity's sake, the general view about the relationship between intentionality and phenomenality that we are going to explore is *the moderate identity view:* intentionality is identical to a kind of phenomenality.

4.1 Phenomenality

On a very plausible understanding of phenomenality, the phenomenality of some mental state is the way that state is for its subject—i.e. how that state is, or what it is like for that subject. We might call the particular phenomenality of some mental state its phenomenal character. Vague though such a definition undoubtedly seems, I take such an understanding of phenomenality to indicate that the phenomenal character of some mental state is in some sense a modifying feature of that state. That is not to say that there is the mental state M and then its phenomenal character P, and that P then modifies M to produce PM. Rather, I take the relevant kind of modification to be more along the lines of how quickly, slowly, briskly, etc. can modify running. It is not as if there can be unmodified runnings— runnings that are neither fast, quick, slow, brisk, etc. Sure, one can run without running quickly, but one cannot run without running in some way. And just as there are not empty runnings that are then modified to become quick runnings, there are not phenomenally empty mental states that are then modified to become phenomenal mental states. The phenomenality of some mental state is not something that can be stripped away from that state to yield the mental state minus its phenomenal character, although we can choose to bracket off, or focus on, some phenomenal features rather than others. In the same way, we can focus on certain properties of the running rather than others. Rather, the phenomenal character of some mental state is one way for that mental state to be, just as quickly, slowly, briskly, etc. are

ways runnings can be. Moreover, just as slowly and quickly are ways
runnings can differ from one another, so too can phenomenalities be ways
mental states can differ from one another. If this is a plausible understanding
of phenomenality, then, given our goal of examining the moderate identity
view above, it follows that the kind of views we are interested in exploring
should construe intentionality as modifying features of some mental states.
But are there any such views?

5 Phenomenal Intentionality Theory

Since the mid to late nineties, a growing number of theorists have become increasingly unsatisfied by the standard picture of the mind that has it divided between the phenomenal and the intentional. Among the first to complain were Searle (1992) and Loar (2003)[57], who, in their own ways, contended that we really have no conception of intentionality/content as divorced from consciousness. Siewert (1998) and Strawson (2010) too were among the view's early advocates, each having made foundational contributions to what is now called Phenomenal Intentionality Theory (hereafter PIT). More recent adherents include Kriegel (2013), Pitt (2009), Horgan and Tienson (2002), Mendelovici (2010), Farkas (2008) and Georgalis (2006).[58]

While a varied bunch, phenomenal intentionalists are generally agreed on several theses (see (Kriegel, 2013) for a more comprehensive list of the central tenets of phenomenal intentionalism). One of the most central of these is that intentionality and phenomenality do not form two separate mental realms, but are instead inseparably intertwined (Horgan & Tienson, 2002). Many also take phenomenality to be the more basic or foundational of the two. This latter point is what sets phenomenal intentionalism apart from representationalism, which also holds an inseparability thesis. Representationalists usually attempt to reduce phenomenality to intentionality, and then account for intentionality in terms of one of the tracking relations vetted in part 1 (see, for instance, Tye, 2002). That being

[57] However, both Searle and Loar seemed to have denied the separatist picture since at least the eighties; see (Searle, 1983) and (Loar, 1987).
[58] However, I think Georgalis would reject the name 'Phenomenal Intentionality'.

said, phenomenal intentionality theory does not require a commitment to the basicness of phenomenality over intentionality. Indeed, several phenomenal intentionalists, including Strawson (2004), Mendelovici (2010) and Pitt (2009), hold theories on which intentionality is identical to phenomenality. And if the two are identical, it does not make much sense to talk about the basicness of one over the other. However, I take the ethos of phenomenal intentionality theory to be different from that of representationalism, even on identity views.[59]

In rough outline, adherents of PIT generally acknowledge a kind of intentionality—phenomenal intentionality—that is in some way grounded in, determined by, or identical to phenomenality (Kriegel, 2013) (Horgan & Tienson, 2002) (Strawson, Mental Reality, 2010) (Siewert C. P., 1998). As a corollary, the intentional content of a phenomenally intentional state is then held to be determined by, or identical, to that state's phenomenal content (Horgan & Tienson, 2002). For the most part, phenomenal intentionalists take phenomenal intentionality to be basic (Kriegel, 2013) or conceptually prior (Searle, 1992) to other kinds of intentionality. Stronger views, such as Strawson's (2010), hold that the only kind of real intentionality is phenomenal intentionality.

5.1 Summary and Look Forward

Here is where things stand. I raised some concerns for some of the most popular relation views of intentionality (ch. 1 – 3)—concerns that I thought warranted exploring non-relational alternatives. From the fact that there appears to be a close connection between phenomenality and intentionality (ch.4), I suggested that we might begin our investigations with this in mind.

---

[59] However, see (Mendelovici, 2010), who disagrees.

What I proposed was that we explore a general view about the nature of the relationship between intentionality and phenomenality—namely the moderate identity view. In the previous section (sec. 4.1), I suggested that phenomenality is plausibly construed as non-relational—a modifying feature of some mental states. Given that the moderate view identifies intentionality with a kind of phenomenality, it would seem to follow that we should be on the lookout for theories on which intentionality is construed as a modifying feature of mental states. On the face of it, phenomenal intentionality theory (or certain versions of it) seems like a promising avenue down which we might find one or more such theories. In the remainder of this chapter, I will examine three versions of PIT that appear to fit the bill.

5.2 A First Approach

In the literature, there are different views about how to construe intentional contents in such a way that intentionality does not turn out to be a relation.[60] One such view, called type psychologism, comes to us from David Pitt (2009). Pitt argues that intentional contents are phenomenal types, of which particular phenomenal intentional episodes [61] are tokens. Intentional episodes are phenomenally constituted such that a particular thought is a token of some phenomenal type that just is that thought's intentional content. To illustrate the point, take the case of a prototypically sensational episode, such as having a pain. For Pitt, one plausible way to understand such states is not as relations to contents, but as being tokens of a particular phenomenal type; namely, the painful type.

---

[60] Hereafter, my use of 'content' should be taken as denoting phenomenal intentional content.

[61] Again, take 'intentional episode' broadly to include thoughts, perceptions, etc.

For states such as pain, such an analysis seems pretty intuitive. The question, however, is how well will it work for prototypically cognitive states, such as believing that it will rain. The concerns here are twofold. First, if we take the claim that intentional contents are phenomenal types, to imply that they have phenomenality, then Pitt's view implies there is something it is like to have the thought that it will rain. Here, as Pitt claims in (2004), having the thought that it will rain can be distinguished from taking a particular propositional attitude towards the proposition 'that it will rain'. The idea is simply that one can entertain the proposition that it will rain without thereby believing, fearing, or remembering it. However, even on this thin construal of what it is to think that it will rain, Pitt's view is that the content 'that it will rain' is a phenomenal type. Hence, there is something it is like to token the content 'that it will rain'.

Though philosophical orthodoxy is, I think, still sceptical of any such phenomenality, there is a wealth of arguments (indeed, an entire volume (Bayne & Montague, Cognitive Phenomenology, 2011)) dedicated to what is called cognitive phenomenality.[62] However, Pitt's account does not merely need there to be cognitive phenomenality; Pitt's account needs said phenomenality to, in some way, be constitutive of thought. If thinking that it will rain is a matter of tokening a phenomenal type that just is that thought's content, then the content 'that it will rain' is a phenomenal type. It is a phenomenal type of thing. In short, Pitt's account actually needs there to be enough phenomenality to the thought that it will rain to distinguish it from all other thoughts. If there is just some generic phenomenality to cognition

---

[62] For an early but especially convincing one, see Strawson (2010).

such that all conscious thoughts have it, then individuating[63] thoughts by their phenomenal character would deliver a pretty short list of thoughts.

Perhaps that was too fast. The idea here is that on a very intuitive way of individuating thoughts—the way I suggested back in part 1—thoughts are individuated by their contents. If contents are phenomenal types, as per Pitt's account, then thoughts are individuated by phenomenal type. That means either that there had better be a sufficiently distinct phenomenal type that can constitute a particular thought's content, or else, that there are not very many types of thought. If the thought that it will rain is the thought it is in virtue of its content, and this content is a phenomenal type, then either every thought is the same (the generic phenomenality view), or else the phenomenality of a particular thought is sufficiently distinct to distinguish it from all other thoughts.

Presumably, Pitt does not want the former. Hence, he needs some story about how the phenomenology of a thought is sufficient to distinguish it from all other thoughts—sufficient, that is, to individuate it. Pitt has such a story, and I think it a good one (Pitt, 2004, pp. 7-14). But rather than rehearsing his arguments, I provide three simpler arguments that I find persuasive, in what I take to be in order of increasing persuasiveness.

5.3 Three Arguments for Distinctiveness:

Assuming that there is cognitive phenomenality—something it is like to think that it will rain—the question before us is whether or not that phenomenality

---

[63] The way I am using it here, to individuate X is to distinguish it from other Xs. Of course, there are better and worse ways to do this: I can stick post-its on MRIs of my brain and therefore individuate my brain states with post-its. But individuating intentional states by their content does not seem like such an objectionable way of individuating them.

is sufficient to individuate that particular thought. Struggle though you might to put into words what exactly it was like to think that it will rain, imagine that things are exactly that way, and ask yourself whether you could be thinking a different thought. That is, if everything is phenomenally like what it is like when you think it will rain, could you possibly be thinking another thought? It is hard to deny that if how things are for you when you think that it will rain are precisely how things are for you now, then you just are thinking that it will rain. The point? If you accept what has so far been said, then you cannot be thinking that it will be sunny, if things are exactly as they are for you when you think that it will rain. Nor could you be thinking that your dogs need supper, that your air conditioning is costing you a fortune, or that there are only 23 species of crocodilian.

Perhaps the above did not convince you; let me try again: Think to yourself that $16 + 32 = 48$. Now, forget trying to describe how it was for you to think that thought, do not bother trying to remember how it was. Simply ask yourself whether your thinking that thought differed in some real, palpable, phenomenally observable way, from a calculator's computing it. The answer is clearly yes. So, though you cannot describe it, there was something it was like for you to think the thought, something you had, but the calculator lacked. Of course, the obvious objection here is that though there might have been some way it was to think that thought (that you had but the calculator lacked), it was a generic what-it-was-likeness—too muted and indistinct to do any work here. But is this right? That is, if how things are, is exactly how they were for you when you thought the thought you did, would you not just be thinking that thought. If what it was like for you to think that $16 + 32 = 48$ is exactly what it is like for you right now, could you possibly be thinking anything else? If the answer is no, then the what-it-is-likeness of thinking

that 16 + 32 = 48, though indescribable, is all that is needed to individuate the thought from all others.

Final attempt: Imagine that you are omniscient and omnipotent with respect to the phenomenality of some subject. Now, imagine that subject thinks 16 + 32 = 48, and you pay close attention to what it was like for him to entertain that thought. Now ask yourself the following: If you arranged things phenomenally for that subject such that he was exactly the same as he was when he thought that 16 + 32 = 48, what would he be thinking?[64] Presumably he would be thinking precisely that 16 + 32 = 48. In which case, fixing the phenomenal character, the how things are, is sufficient for fixing the content of 16 + 32 = 48-thoughts.

The point is this: If these arguments are right, then cognitive phenomenality is indeed sufficient to deliver intentional individuation, and Pitt's account is viable up to this point.

5.4 Bonafide Thoughts and the Propositional Attitudes

Earlier, I said that there were two possible concerns with Pitt's account. The first was that Pitt needed the phenomenal character of thoughts to be sufficient to individuate them. I have considered what I take to be plausible reasons for thinking this might be possible. The second concern has to do

---

[64] The reason why I like this kind of argument is that it does not rely on actually describing, in distinct terms, the phenomenal character of the various thoughts. Indeed, it need not even require that the phenomenal character be, in principle, describable—other than, of course, being describable as the phenomenal character of thinking that 2 + 2 = 4, or that it will rain.

with thoughts construed not thinly as mere occurrences of a particular content, but as full on propositional attitude states. The problem, that is, is that when the thought that it will rain occurs, it rarely does so as a mere entertaining. When the thought that it will rain occurs to me, it typically does so in the context of my believing, fearing, remembering or desiring that it will rain. Remember that on Pitt's account, contents are phenomenal types; so if the content of my belief that it will rain is that it will rain, and the content of my fear that it will rain is that it will rain, then it looks like we have two thoughts that token the same phenomenal type. But then what resources does Pitt's account have to distinguish the two thoughts? I assume that we would all like the desire that it will rain to count as a different thought than the fear that it will. But if token thoughts are individuated by their contents, which are phenomenal types, then what exactly distinguishes the desire that it will rain from the fear that it will?

One option here is to recognize that the above is only a problem if you accept that thoughts are *wholly* individuated by their contents. On one traditional view, thoughts (in this rather thicker sense) are individuated by their attitude + their content. So though the desire that it will rain and the fear that it will have the same content, the two take a different attitude toward said content and so are different thoughts. One move open to Pitt would be to tell some story about how each attitude type has a corresponding phenomenality such that this attitude + content schema could deliver the individuation of bonafide propositional attitude thoughts. The idea here would be that belief states have the phenomenal character of believing, which, in conjunction with the phenomenal character of the content that it will rain, would yield a unique phenomenally individuated propositional attitude state (see Horgan & Tienson, 2002 for such an approach).

Another option here is to reject the attitude/content distinction and hold that propositional attitudes are part of the content of intentional states. On this line the difference between believing that it will rain and fearing that it will is a difference in content.[65] Personally, I think this is the better option, especially for a view like Pitt's. After all, there does indeed seem to be something it is like to believe something that is different from what it is like to fear it. Taking such attitudes to be part of the content of intentional states explains why believing that it will rain is a different thought than desiring that it will. Each have different contents; each tokens a different phenomenal type that includes the attitude type.

5.5 Two More Approaches

Two other phenomenal intentionalists have provided similar alternative views. According to Kriegel's adverbialism (2011), and Mendelovici's aspect view (MS.), intentional contents are not phenomenal types, but second order properties of intentional properties. On Kriegel's adverbialism, contents are taken to be adverbial modifications of intentional states. Where this view differs from the adverbialism prominent in the writings of theorists such as Ducasse (1942), and Chisholm (1957), is in its scope, as well as its focus. For the early adverbialists, the goal was to provide a plausible alternative to sense data theory, one that eschewed mind-dependent sense data for adverbial translations of our perceptual talk. Rather than 'I see red' meaning that the speaker is in direct contact with a red sense datum, it means instead that the speaker perceives redly, or red-wise. On Kriegel's adverbialism, however, it is not just perceptual experiences that are adverbially construed, but all intentionality. Also, the focus is not on providing adverbial

---

[65] In his doctoral thesis, *A Pure Representationalist Account of the Attitudes*, Steve Pearce develops such a view in far greater detail than I do here.

translations, but on giving an account of being intentionally directed that is both phenomenally based, and able to cope with various issues arising from the relation view (see Mendelovici, MS, sec. 8.4.2).[66]

5.6 Adverbialism and The Aspect View

For Kriegel then, having a thought about a dog is a matter of instantiating an intentional property, the property of being intentionally directed, that has certain modifying properties that may be glossed as dog-wise. Again, these properties of intentional properties are phenomenal properties: If the adverbialist construal of contents takes them to be ways intentional properties are, and phenomenal characters are ways phenomenally intentional states are, we get a rather nice phenomenal construal of contents.

On one version of Mendelovici's aspect view (MS), contents also turn out to be second order properties of intentional properties. Where this version of her account differs from Kriegel's is in its understanding of the second order properties in question. On Mendelovici's view, the relevant second order properties are construed more along the lines of the second order properties of color; namely hue, saturation, etc. (Mendelovici, MS). Of course, the second order properties in question do much of the same modifying work as Kriegel's adverbial properties—i.e. they are ways intentional properties are, or can be—but, as with the case of hue and saturation, thinking of them in adverbial terms is not the most natural way to conceive of them. We do not, for instance, say that red_2 has a 2-ly hue.

A final important feature of both Kriegel's and Mendelovici's respective accounts is that the relevant second order properties do not compose—as

---

[66] See Part 1 for some such issues.

distinct constituents—the intentional properties they modify. While Kriegel does not argue for it explicitly, adverbial properties are not distinct constituents of what they modify. Quickly, for instance, does not compose runnings. On the other hand, and in keeping with the color analogy, Mendelovici explicitly says that just as we need not say that properties such as hue and saturation are distinct components that literally compose the color properties they modify, so too for second order intentional properties (Mendelovici, MS, ch. 8). This is important because the whole idea of taking contents to be second order properties is to avoid giving them "independent existence" status—and therefore status as things we are related to via intentionality. Taking the relevant properties to be of the first order would be to endorse the sense data theory.

To reiterate: On the adverbial/aspect view, intentional contents are second order phenomenal properties of intentional properties; namely they are ways for those intentional properties to be rather than distinct constituents of said properties. This differs from Pitt's view that we looked at earlier. On his view, contents are phenomenal types of which particular intentional episodes are tokens. As we saw, the problem for Pitt's view was that it had the rather troublesome consequence that I cannot desire what you fear. This resulted from the combination of several points. The first was that Pitt's account is a non-relational account, and therefore rejects the relational view of propositional attitudes. The second was that in order to avoid positing things such as empty believings, Pitt's account would have to reject the attitude/content distinction, and say that the attitudes are part of the content. But since contents are phenomenal types, on Pitt's view, this would mean that my desire that it will rain is of a different type than your fear that it will. And this seemed to imply that my desire and your fear share no content.

Does adverbialism or the aspect view fare any better here? Indeed, I think both do. The reason why I desire what you fear is that our respective intentional states share some second order properties; namely, rain-wise (as per Kriegel), or of-rain-ness (as per Mendelovici). Of course, where Kriegel and Mendelovici stand on rejecting the attitude/content distinction I do not know. Perhaps one or either would say that the attitudes are the first order properties that the second order properties modify. My believing it is raining is a matter of my having a belief state that occurs rain-wise, or that has of-rain-ness. And just as there are no runnings that are not quick, slow, etc. runnings, so too are there no beliefs that are not rain-wise, or of-redness, etc. This seems like a win/win: We need not reject the age old attitude/content distinction, and we get a non-relational view on which distinct subjects can be said to share some contentful commonality.

There is, however, a worry that emerges for the views under consideration. The theorists we are considering take themselves to be espousing phenomenal intentionality theories—theories that, we might say, tie intentional content to phenomenal character in an essential way. If we accept that intentional content is a kind of phenomenal character, there ought not to be any phenomenal difference without a corresponding contentful difference. But how it is for me to desire that it rain is quite different from how it is for me to fear it. That is, there is a phenomenal difference between the two. But the content, on a view that accepts the attitude/content distinction is the same. In particular, the second order properties of rain-wise, or of-rain-ness, seem to be the same in both the case of desiring that it rain and fearing that it will.

Of course, the solution here is obvious. The views we are considering ought to reject the attitude/content distinction. They ought, that is, to take propositional attitudes to be contents like any other; namely, second order properties of intentional properties. In this case, desiring that it will rain is more accurately understood as being intentionally directed in a desiring, rain-wise way. And fearing that it will rain is a matter of being intentionally directed in a fearing, rain-wise way.

5.7 Summary

I began part 2 by suggesting that we explore views on which the intentional and the phenomenal are closely connected. In particular, I said that I wanted to examine what I called the moderate identity view. This led me to phenomenal intentionality theory, and I examined three different theories that I take to be consistent with this view. I began with Pitt's type-psychologism, but I concluded that Pitt's account had some consequences that I find unpalatable. I then looked at two similar alternative accounts and argued that the unpalatable consequences of Pitt's view do not arise for them, provided some conceptual fine-tuning. What is more, the latter two views—Kriegel's adverbialism, and Mendelovici's aspect view—accord well with what I called a plausible understanding of phenomenality; namely, that phenomenality is a modifying feature of mental states. Both adverbialism and the aspect view take intentional contents to be modifying in this way.

Hereafter, I refer to both Kriegel's adverbialism and Mendelovici's aspect view as versions of *modificationism*. Modificationism combines the moderate identity view (call this moderate modificationism) with the thesis that phenomenal character is a way for a mental state to be—a modification of that state. Hence, intentional content is a way for a mental state to be—a

modification of that state. In the next chapter, I will highlight some challenges that modificationism faces.

6. Challenges to Modificationism

In this chapter, I highlight some challenges to modificationism. The goal of this chapter is not to provide decisive arguments either in favour of or against modificationism, but rather to see how far it can be pushed, with the purpose of identifying its weaknesses, in order to see how a sympathizer might reasonably seek to overcome those weaknesses.

The first family of challenges for modificationism arises from cases where the intentionality of an experience appears to outstrip its phenomenality. Such worries generally emerge in discussions of cognitive phenomenology. For instance, is there anything it is like to think that 2 + 2 = 4?

A second set of concerns surrounds the fact that modificationism is a version of phenomenal intentionality theory, and phenomenal intentionality is…well…phenomenal. Assuming that phenomenal intentionalists are agreed that it is phenomenal intentionality that is, in some sense, the most basic form of intentionality, the view seems incompatible with the very possibility of unconscious intentionality. More specifically, phenomenal intentionality theory, and therefore modificationism, seems unable to accommodate unconscious, yet seemingly contentful, mental states. Worries here differ according to the nature of the relevant unconscious states. Standing states such as the standing belief that the Eiffel Tower is in France pose one set of difficulties, while states of the early visual processing system pose another. Concerns also abound about the content of subconscious states—the ones that, according to some, figure crucially in how we behave.

A third concern that is specific to modificationism has to do with the compositional structure of thought. In Jackson (1977, pp. 64 – 72), some serious challenges are lodged against adverbial theories of perception. As a

form of adverbialism, modificationism would appear vulnerable to these same challenges.

Finally, there is the rather large concern that the views we are considering leave us all trapped inside our own heads. Deprived of the relation that was once thought to constitute intentionality, one might worry that our cognitive contact with the world—the contact secured by that very relation—has gone the way of the dodo. In effect, it looks like the present views do not have the resources to connect us, our minds, our thoughts, etc. to the world of ordinary objects that we inhabit.

In summary, we have four sets of concerns that need to be addressed:

1) There seem to be cases where the intentionality of some experience outstrips its phenomenality.
2) Phenomenal intentionality theory (and therefore modificationism) appears inconsistent with the existence of unconscious, yet contentful, mental states states.
3) Modificationism seems inconsistent with the idea that thoughts are structured in a particular way.
4) Modificationism appears unable to deliver cognitive contact.

In the remainder of part 2, I will try to enumerate some options a sympathizer might have with respect to concerns 1) – 3). 4) Will occupy the entirety of part 3, so I will not address it here.

97

6.1 Cognitive Phenomenology

As mentioned, the concern here is that certain intentional states have an intentionality that outstrips their phenomenality. In effect, this concern is about both the existence and the distinctiveness of cognitive phenomenology. This is a popular topic as of late, garnering an entire anthology (see Bayne & Montague, 2011). The question is whether there is anything it is like to undergo prototypically cognitive states such as believing that it will rain, or wondering what the square root of pi is.

In Pitt (2004, p. 2) an initial—albeit trivial (or so Pitt claims)—argument for the existence of cognitive phenomenology runs thus:

1) If a mental state is conscious, then it has phenomenal properties.
2) Conscious thoughts are conscious mental states; therefore,
3) Conscious thoughts have phenomenal properties.

The argument looks to be valid, and therefore to argue against it will involve arguing against its premises. While most will admit that there is often a what-it-is-like-ness that accompanies such states, those sceptical of a phenomenality distinctive of cognition typically try to account for this in terms of sensory phenomenality (see Prinz, 2011). The idea is to deny 1) by arguing that just because a conscious mental state occurs contemporaneously with certain phenomenal properties does not mean the latter are properties of the former. I can, for instance, be doing long division while listening to Mozart, but the phenomenal character of that episode is not distinctive of long division, but of listening to Mozart. On this view, in other words, there can be a phenomenal character *with* cognition, but there is no phenomenal character *of* cognition.

98

It follows that many of the arguments for cognitive phenomenology focus on cases where there is a phenomenal difference that cannot be accounted for in terms of sensory phenomenology. In Galen Strawson's famous example, two subjects—one of whom does not speak French—listen to a French news broadcast (Strawson, 2010, pp. 5 -13). By hypothesis, both have identical perceptual/sensory experiences; yet their phenomenal experiences differ. According to Strawson, the difference is a difference in the phenomenality of understanding, or understanding experience.   Since understanding is precisely the kind of cognitive state at issue, there must be something it is like to think.   Other such examples include the what-it-is-like of having something on the tip of one's tongue, the experience of grasping what 'dogs dogs dog dog dogs' (Horgan & Tienson, 2002, p. 523) means, and of suddenly remembering what you forgot.

Such examples at least establish a plausible case for the existence of cognitive phenomenology.   However, something more is required before we can conclude that a cognitive state's phenomenality can be sufficient for its intentionality.   More precisely, what modificationism requires is some account of how some mental state's intentional content can be constituted or individuated by its phenomenal character.

That being said, Strawson's example might be construed (or suitably tweaked) to establish as much. Kriegel, for instance, suggests a case where:

> "two languages are so similar graphically and phonetically that the very same passage can express a news report about a faraway war in one of them, and a children's bedtime story in the other. We can envisage two subjects listening to a reading of the passage and each understanding it in a different language.   Here there would be an overall experiential

difference that is best explained by supposing that one subject's cognitive experience had one intentional content while the other's had another intentional content" (Kriegel, 2011, p. 49).

Pitt, on the other hand, takes a different approach—taking as his starting point the fact that we can consciously, introspectively and non-inferentially know about our phenomenal states, and that it is only conscious experiences that are so knowable. But this is precisely the kind of knowledge we have of our thoughts, and "it would be impossible to introspectively distinguish conscious thoughts with respect to their content, if there weren't something it is like to think them" (Pitt, 2004):

> Normally—that is, barring confusion, inattention, impaired functioning, and the like—one is able, consciously, introspectively and non-inferentially (henceforth, "Immediately") to do three distinct (but closely related things): (a) to distinguish one's occurrent conscious thoughts from one's other occurrent conscious mental states; (b) to distinguish one's occurent conscious thoughts each from the others; and (c) to identify each of one's occurrent conscious thoughts as the thought it is (i.e., as having the *content* it does). But…one would not be able to do these things unless each (type of) occurrent conscious thought had a phenomenology that is (1) different from that of any other type of conscious mental state (proprietary), (2) different from any other type of conscious thought (distinct), and (3) constitutive of its (representational) content (individuative) (Pitt, 2004, pp. 7-8).

To be sure, this kind of self-knowledge argument is open to several possible counter-arguments (see for instance Tye & Wright, 2011), some of which are addressed in (Pitt, 2011). For instance, one might flat out deny that we have

any such knowledge of our mental states, be they intentional, phenomenal, or phenomenal intentional. However, the point here is not to provide an extended argument in favour of cognitive phenomenology, but to lay out several options a modificationist might reasonably pursue. Pitt's view, if successful, would certainly give the modificationist what she needs.

From Horgan and Graham (2012) and Horgan and Tienson (2002), we find another possible avenue for securing the kind of cognitive phenomenology the modificationist needs. According to them, it is the phenomenal character of thinking that rabbits are furry that makes the content of that thought *that rabbits are furry*, rather than *that the set of undetached rabbit parts is*. Horgan, Graham and Tienson seem to think it obvious that there is something it is like to think that rabbits are furry, and that that is quite different from what it is like to think that the set of undetached rabbit parts is. If they are correct, then phenomenality delivers determinate content.

For my part, I too take it that what is at issue between friends and foes of cognitive phenomenology is not the existence of certain phenomenal features during episodes of cognition. All sides grant that I can do long division while enjoying the phenomenal experience of listening to Mozart. What is in dispute is whether there is a purely cognitive phenomenality—phenomenality that is proper to cognition and not reducible to sensory (or otherwise) phenomenality, and that may or may not occur simultaneously with cognition. Siewert (2011, pp. 262 - 267) offers a useful way to understand this point by dividing phenomenality into derived and non-derived. The phenomenality of doing long division while listening to Mozart is different from doing it while listening to your favourite Punk band. Notice too that the difference need not just be in the phenomenality of the aural perceptual experience: it really seems like the phenomenality of the aural

perceptual experience somehow bleeds through into the task of doing long division, so that the cognitive exercise of doing long division is coloured by the former. But this would still not count as cognitive phenomenality proper because the phenomenality of the cognitive task is derived from that of the sensory experience.

Siewert uses other examples to try and bring out a kind of non-derivative cognitive phenomenality, and goes on to argue that this phenomenality is constitutive of content (2011, pp. 262 - 267). However, even if Siewert is mistaken and all phenomenality does in fact derive from the sensory, I am not sure that that constitutes a huge problem. For argument's sake, let me grant that all phenomenality is in some way sensory, so that any phenomenality that accompanies cognition is derived in Siewert's sense. Using Siewert's example of reading with comprehension versus reading without comprehension, we get a case of phenomenal contrast. There is something it is like to read with comprehension—something different from what it is like not to comprehend what you are reading. Let us assume, on behalf of the foes of cognitive phenomenology, that the specific what it is likeness of reading with comprehension derives from the sensory phenomenology that accompanies your internal monologue in which the words are present. The question is whether the derivative nature of the phenomenality presents a problem. What modificationism requires is that the phenomenal character of some cognitive episode can constitute, or in some way individuate, the intentional content of said episode.

Imagine removing whatever phenomenality arises as a result of your internal monologue that occurs when you read with comprehension. That is, imagine that you no longer have any sensory phenomenality arising from your internal monologue. In such a case it is not unreasonable to think that you

also lose comprehension. Bring the phenomenality of one's internal monologue back, and with it comes comprehension. This suggests that phenomenality—whether it be derived from the sensory or not—is in some sense constitutive of reading with comprehension. So even if we accept that all phenomenality is in some sense sensory—or derived from the sensory— that does not necessarily pose a serious threat to modificationism. What the modificationist would require here would be an argument to the effect that conscious thought is sensory in the relevant sense. Prinz (2007) offers something like this kind of view. According to Prinz, all mental representations are perceptual in nature, and therefore have perceptual phenomenology (Prinz, 2007, p. 348). Perhaps this perceptual phenomenology derives from one's internal verbal narrative, or mental imagery, or some such. What is important is that its being sensory (perceptual) does not entail that it cannot constitute or individuate the content of the relevant thought.

In summary, our stated purpose in this section was to investigate whether modificationism faced some insurmountable challenges stemming from its seeming reliance on both the existence, and content determining powers, of cognitive phenomenology. While I admit there is a great deal more work required here to vindicate modificationism on the cognitive phenomenology front, I also submit that the challenges modificationism faces here are not insurmountable. Indeed, there appears to be promising progress here.

6.2 Unconscious Content

If you recall, modificationism is committed to the view that original intentionality is a kind of phenomenality—that (original) intentional content is phenomenal. Non-conscious states that have intentional content thus

stand as immediate counter-examples to modificationism. It would be extremely convenient if non-conscious states formed a well-behaved kind such that they could all be addressed together. Unfortunately they do not. There are three broad categories of non-conscious states that could pose challenges for the modificationist: 1) states of the deep sub-conscious, such as those that might make one want to kill his father and marry his mother; 2) standing states, such as one's non-occurrent belief that the Eiffel tower is in France; 3) sub-doxastic/sub-personal states, such that those of early visual processing.

In what follows, I list several proposals for handling (some) non-conscious states (or others). My suggestion will be that given the disparity in kind of non-conscious states, some combination of the following views is the best way to handle the cases.

Searle (1992, pp. 155 – 162) argues for what he calls the *connection principle*. The view is that non-conscious states have the intentional content they do in virtue of the phenomenal character they would have were they conscious. This suggests that the relevant kind of unconscious states are at least *potentially* conscious, and that their intentional content can only be determined relative to their connection to consciousness. In a similar vein, Mendelovici suggests a kind of dispositionalism about some non-conscious states. In particular, one's standing non-occurrent belief that the Eiffel Tower is in France is simply one's disposition to occurently believe that the Eiffel Tower is in France (Mendelovici, MS).

Horgan and Graham (2012, p. 341) offer a position dubbed *inferentialism* by Kriegel (2011, p. 194). The idea here is the intentional content of subconscious states is derived from:

"...their role in an overall assignment of contents to actual and potential states of the cognitive system under which: (i) all phenomenally conscious states are assigned the content that they inherently possess by virtue of their intrinsic phenomenal character, and (ii) all other states in the system are assigned contents in such a way that the overall content-assignment exhibits and acceptably high degree of internal rational coherence" (Horgan & Graham, 2012, p. 341).

This view emerges within the framework of arguing that only phenomenal intentionality (intentionality constituted by phenomenal character) has determinate content. And it is the determinateness of this phenomenal intentional content that is supposed to do the work of ensuring that a unique content is assigned to the non-conscious states. The idea is that the determinate phenomenal intentional contents act as a network of "anchor points" sufficient to *infer* a unique content to each subconscious state. Subconscious states derive their content by inference—inference based on the determinateness of phenomenally intentional states and how the relevant subconscious states interact with them in the "cognitive architecture of competent human cognizers" (Horgan & Graham, 2012, p. 341).

Kriegel's view, *interpretivism*, is that non-conscious states derive their intentional content by interpretation (Kriegel, 2011, pp. 200-218). A non-conscious state, S, has the intentional content that P, if the best interpretation of the cognitive system to which S belongs would assign P to S. Kriegel credits this view's starting place to Dennett (Kriegel, 2011, p.201). Indeed, interpretivism in this sense involves taking the *intentional stance* towards the sub-system that constitutes a cognitive system's non-conscious states.

Another interesting suggestion from Mendelovici (2010) concerns sub-personal/sub-doxastic states. Her view is that though such states are not intentional in the same sense that conscious states are, nothing stands in the way of attributing them a kind of informational/computational content (Mendelovici, MS)—call this view *informationalism*. This content might be attributed to them according to certain causal co-variational relations they bear to things in the cognitive system's environment, or by some other means. This does not contradict what was discussed back in Part 1. There, the challenges raised against causal covariational theories were based on the difficulty of specifying which causal/covariational relation was the one that determines *the* content. To repeat, the informationalist holds that there need not be any determinate fact of the matter about what the informational content of the relevant state is. States of the early visual processing system might causally covary with several things in a long causal chain, and the present proposal is that they can be about one, two, or all of these things based on what information we are interested in. As Mendelovici suggests, informationalism has the benefit of "freeing-up" informational content from certain constraints we place on bonafide intentional content (Mendelovici, MS). Consider the rings on a tree. The rings on a tree can carry information about the age of the tree, the chemistry of the atmosphere at a given point in history, the life cycle of certain boring insects, and so on. The virtue of taking sub-doxastic/sub-personal states as having this kind of informational content, is that there need not be any determinate fact of the matter about what they represent/are about: they can be about one, all, or none of these things depending on our interests. On the other hand, phenomenal intentional content does have this kind of constraint (Mendelovici, MS).

As a reminder, there were three kinds of non-conscious states I mentioned earlier: 1) states of the deep sub-conscious, such as those that might make

106

one want to kill his father and marry his mother (subconscious states hereafter); 2) standing states, such as one's non-occurrent belief that the Eiffel tower is in France; 3) sub-doxastic/sub-personal states, such that those of early visual processing. Each poses a prima facie challenge for modificationism—the view that intentionality is a kind of phenomenality—since non-conscious states are non-phenomenal. As I said earlier, it would be very surprising if a single theory of the non-conscious could account for such disparate kinds of non-conscious states.

Again, my goal here is to see whether modificationism faces any insurmountable challenges arising from non-conscious states. I have listed several theoretical options found in the literature, and suggested which ones are open to the modificationist. To repeat, I am agnostic about which approach best accounts for subconscious states of the sort posited by Freud. As for standing states, I favour dispositionalism for reasons of content determinacy. I am inclined to attribute a level of determinacy to my standing non-occurrent states, such as my belief that the Eiffel Tower is in France—a level of determinacy that that I am not inclined to attribute to sub-doxastic/sub-personal states. Dispositionalism helps explain why I have this inclination. My inclination to attribute determinacy to my non-occurrent standing belief that the Eiffel Tower is in France arises because that very non-occurrent belief just is my disposition to have the occurrent belief with that determinate content.

For sub-doxastic/sub-personal states, I am inclined towards informationalism. My reason for doing so is that informationalism seems to be the most compatible with the edicts of cognitive scientific research. Informationalism holds that there are all sorts of informational states that are not accessible to consciousness, but are nonetheless contentful. The

informationalist simply holds that the kind of content such states have is different in kind from phenomenal intentional content.

Before moving on, I would like to note that every view we examined—the connection principle, dispositionalism, inferentialism, interpretivism and informationalism—made the content we attribute to the various non-conscious states somehow derived. Non-conscious states derive their content from the occurrent states they are disposed to be (Searle, Mendelovici), or from an inference based on the determinateness of phenomenally intentional states and how the relevant subconscious states interact with them in the cognitive architecture of competent human cognizers (Horgan et al.), or from the interpretation given them by an ideal interpreter (Kriegel), or from our interests (Mendelovici). In short, if we insist that these disparate non-conscious states are intentional, they are all cases of derived intentionality: they are about what they are about in virtue of intentional states distinct from themselves.

6.3 Structural Modifications

Finally, a family of problems—originating from Jackson (1977, pp. 63-72)—emerges for the kind of view we are exploring.   Jackson's concerns are directed towards the adverbial theory of perception championed by theorists such as Chisholm (1957). However, Jackson's worries would seem to apply equally to the views we are examining. First, in the case of more complex adverbial states such as thinking redly-squarely, it is unclear what the relevant adverbs are modifying. For instance, does redly modify squarely or vice versa?

Consider Jackson's example:

1) He speaks impressively.

2) He speaks impressively quickly.

In 1) 'impressively' modifies the speaking, but in 2) it modifies the other adverb. But in the case of thinking squarely, versus thinking redly squarely, it is not clear that there is any principled way of determining which adverb modifies what.

On behalf of the adverbialist, one is tempted to say that neither adverb modifies the other, but both modify the thinking. When I think redly squarely, I am thinking redly and squarely. Compare: When he speaks, he speaks impressively and quickly. Several issues emerge for this way of understanding the adverbial properties. First, when I think white-unicorn-ly, it seems like the 'white-ly' is somehow connected to the 'unicorn-ly'. This is why I can easily entertain white-unicorn-ly thoughts, as opposed to tree-car-ly, or table-chair-ly thoughts. But how can the present view—on which the white-ly and unicorn-ly both modify the thinking—make sense of this?

Second and relatedly, imagine thinking about a white unicorn and a hairy Bigfoot. On one adverbial construal, we are to understand this as thinking white-ly and unicorn-ly and hair-ly and bigfoot-ly. But this is precisely how someone would be thinking were they thinking about a hairy unicorn and a white Bigfoot. The adverbialist cannot distinguish thinking about a white unicorn and a hairy Bigfoot from thinking about a hairy unicorn and a white Bigfoot (Jackson, 1977, p. 64).

What to do? One move is to fall back on the original suggestion that when one thinks about a white unicorn and a hairy Bigfoot, one is thinking white-unicorn-ly and hairy-bigfoot-ly. The problem here is that while it can be

inferred that I am thinking about unicorns from the fact that I am thinking about white unicorns, the same does not appear to be true of my white-unicorn-ly thoughts (Jackson, 1977, pp. 69 – 72). In a similar vein, it seems that when I think about a white unicorn, and then a grey unicorn, my thoughts somehow overlap, or have something in common—namely, they are both about unicorns (Jackson, 1977, p 67). However, the adverbial construal cannot accommodate this: 'unicorn' is no more a part of 'white-unicorn-ly' than 'straw' is of 'strawberry'. The problem has to do with the structure of thought: Something about the *structure* of the thought licenses the inference from 'I am thinking about white unicorns' to 'I am thinking about unicorns'. Likewise, something about the structure of the two thoughts—the one about white unicorns and the one about grey ones—makes them similar: Each appears to have 'unicorn' as a constituent.

What this comes down to is that modificationism seems unable to accommodate the fact that our thoughts are structured—being *composed* of simpler parts in a *systematic* way so as to *produce* a complex whole. In short, the objection is that modificationism cannot accommodate compositionality.

One move for the modificationist here is to bite the bullet and deny that thoughts have compositionality, or have it only derivatively. One might, for instance, argue that though thoughts do not have *inherent* compositionality, they somehow manage to derive a kind of pseudo-structure from the structure of language. The idea here would be that language somehow contributes to the expressive power of thought so as to imbue thought with the requisite kind of structure. Prima facie, this seems at odds with the idea that the intentionality of words derives from that of thoughts, but it need not be. For instance, I can use sticks and stones to build a complicated machine whose purpose is to manufacture hammers and nails, with which I can build

an even more complicated machine whose purpose is to build air compressors and pneumatic nail guns, and so on. Thoughts (or concepts) would be the sticks and stones, and language would be like the machine/s. From fairly rudimentary (non-compositional) elements, I can build a more complicated (combinatorial) system, which then serves to shape, transform and empower the rudimentary elements into exponentially more powerful tools. I think it is an interesting view, and one that a modificationist might reasonably pursue.

On the other hand, it might be argued that modificationism can indeed deliver the relevant kind of structure. A sketch of such an argument can be found in Kriegel (2011, pp. 161-163). The idea is that perhaps the determinate/determinable structure might help the adverbialist answer some of Jackson's challenges. Very roughly, determinables are general ways things can be, and determinates are more particular ways things can be. Many things in the world exhibit this structure. Red is a determinable of which crimson is a determinate. Maple is a determinable of which Japanese Maple is a determinate. Stephen Yablo has capitalized on this idea to give an account of mental causation (Yablo, Thoughts: Philosophical Papers Volume I, 2008), and answer questions about essence and identity (Yablo, 2010).

With respect to Jackson's challenges, the first was that in cases such as thinking redly squarely, it is not clear which adverb modifies the other. This prompted us to amend the adverbial construal so as to make conspicuous that the redly and squarely modify the thinking, not each other. But this was problematic since, in cases like thinking white-ly unicorn-ly, the 'whitely' seems to be connected to the 'unicorn-ly'. Again, this is merely a sketch of a possible solution, so bear with me. First let us take thinking somehow (being directed somehow) as the ultimate determinable, of which thinking unicorn-

ly is a determinate. Thinking unicorn-ly is also a determinable, of which thinking white-unicorn-ly is a determinate, that is, a particular manner or way of thinking. Let us also assume a level of bottom up transitivity so that if C is a determinate, or modifier of B, and B is a determinate/modifier of A, then C is also a modifier of A, by modifying what modifies A—namely B. With this rudimentary structure in place, we can answer Jackson's first challenge: When I think about white unicorns, I am thinking white-unicorn-ly, which is a determinate of thinking unicorn-ly, which is itself a determinate of thinking. The 'white-ly' modifies the thinking by modifying the thinking unicorn-ly.

But what about thinking red-squarely; which is the determinate or the determinable here? That is, in the case of white unicorns, the order of determinable/determinate is obvious, but it is not so obvious in the red square case. Perhaps this is true, but this does not seem like a problem particular to modificationism. One can imagine asking someone to think of a red square and then asking her what it is she is thinking of in the first place, the redness or the squareness. Sometimes there will be an answer to this kind of question, and other times there will not. In cases where there is such an answer, the present suggestion is that modificationism can appeal to the determinate/determinable structure to accord with it. I can think of a white egg, and I can think of painting my walls egg white. According to modificationism, in the former case, I am thinking white-egg-ly, and in the later I am thinking egg-whitely—where the difference is a difference in the determinate/determinable structure.

The second concern we looked at was that modificationism could not make sense of the fact that, when we think about white unicorns and hairy Bigfoots, it is precisely that and not hairy unicorns and white Bigfoots about which we think. I suggested that the modificationism could avoid this

112

particular concern by understanding the relevant thought as a white-unicorn-ly and hairy-bigfoot-ly thought. The problem with this move was that it seemed to block certain intuitive inferences. From the fact that I am thinking about a white unicorn, it follows that I am thinking about a unicorn. This is because 'unicorn' is a constituent of 'white unicorn'. But the inference does not appear to go through for thinking white-unicorn-ly. This is because 'unicorn' is not a proper constituent of white-unicorn-ly: the latter is syntactically simple. Notice, however, that though the fact that I kicked a strawberry does not entail that I kicked a straw, it does entail that I kicked a berry (Kriegel, 2011, p. 162). This is because a strawberry is a determinate of the determinable berry. If we assume that thinking white-unicorn-ly is a determinate of the determinable unicorn-ly, then the inference that I am thinking unicorn-ly from the fact that I am thinking white-unicorn-ly seems go through.

Admittedly, this is only slightly more of a sketch than is found in (Kriegel 2011), but I think it is a promising one. At some level the demand for compositionality is a demand for structure. To require that thoughts have compositionality is to demand that they be structured in a particular way, or perhaps, that they be capable of structure. And while the determinate/determinable structure is different from compositional structure, it is still structure—structure that seems helpful in answering some of the most famous objections to adverbialism (and hence modificationism). I think the question of how much work this determinate/determinable structure can do for the modificationism is both fascinating and a worthy candidate for another monograph. Stephen Yablo, for instance, argues that the determinate/determinable structure can help us get some traction on several important issues, including mental causation (Yablo, Thoughts: Philosophical Papers Volume I, 2008).

I conclude part 2 with a summary and look forward. Having expressed some dissatisfaction with relation views of intentionality (Part 1), I said part 2 was going to be dedicated to exploring non-relational alternatives. I suggested that we begin our exploration with views on which the intentional and the phenomenal are closely connected. In particular, I said that I wanted to examine what I called the moderate identity view: intentionality is a kind of phenomenality. This led me to phenomenal intentionality theory, and I examined three different theories that I took to be consistent with moderate identity. Dubbing any account that combines the moderate identity view with the thesis that phenomenality is a modifying feature of mental states 'moderate modificationism', I suggested that moderate modificationism appears to fit the bill of being a non-relational account of intentionality. That being said, I explained that several challenges arise for this kind of view, three of which were the problem of cognitive phenomenology, the problem of the unconscious, and the challenge that thoughts are structured. I discussed several options that a modificationist might reasonably pursue in attempting to meet these challenges, and explained which, if any, I thought the most promising. I concluded that the challenges considered are not insurmountable. However, there remains a final difficulty for modificationism, one that was not addressed in part 2—the problem of cognitive contact. To my mind, it is the most pressing of all concerns, since, as will be explained, modificationism seems primed to leave us all trapped inside our own heads. It is to this that I now turn.

Part 3: The Problem of Cognitive Contact

7. Introduction to the Problem

114

So far, we have examined some difficulties with the relational view of intentionality, and have been looking into non-relational alternatives. In particular, the view we have been examining, modificationism, takes intentionality to be a species of phenomenality. On this view, the phenomenal-intentional content of a phenomenally intentional state is an intrinsic modifying property of that state. In Part 2, I considered several purported difficulties with this view, but argued that they could be overcome. In this final section, I want to address another more serious problem with phenomenal intentionality theory in general, and modificationism in particular. The problem is this: it seems like the whole point of having intentionality is that this phenomenon puts us in cognitive contact with the world outside our skins. Introspectively our intentional states appear to accomplish precisely this. But how is this possible on the non-relational account I have presented?

Put another way: It looks as though, almost by definition, making cognitive contact with something involves being related to that thing. If the ability to make cognitive contact with the world is a desideratum on a theory of intentionality, and if making cognitive contact with something involves being related to that thing, then it looks as though any view that denies that intentionality is a relation will, by that very denial, make cognitive contact impossible. Of course, more will need to be said about what exactly cognitive contact is supposed to amount to, but it is not difficult to see how cognitive contact poses a prima facie problem for the views we have been looking at. Indeed, the problem of cognitive contact would appear to make trouble for any account that combines the thesis that intentionality is entirely a matter of phenomenality with the view that phenomenality is an intrinsic, non-relational phenomenon. This is a hard problem, and one that has not been

given enough attention within what Kriegel calls the Phenomenal Intentionality Research Program (2013).[67]

Perhaps this lack attention is due to the fact that the problem of cognitive contact is considered to be a problem for perception and perception alone. After all, it is in perception that our cognitive contact with the world is most apparent. When I have a perceptual experience as of a Japanese maple in leaf, my perceptual experience seems to put me in cognitive contact with the maple. However, the problem of cognitive contact cuts much deeper. When, for instance, I consider rescuing another dog, it seems as though my thoughts are connected to flesh and blood dogs.

To put things in a way more consonant with what has so far been said: the non-relational way one is intentionally directed does not appear to be a good candidate for securing cognitive contact. So, what to do? Perhaps a good starting place is to concretize the problem a bit:

1) We are in cognitive contact with the world outside our skins.
2) It looks as though this contact is established by intentionality.
3) Being in cognitive contact with X is, or is partly constituted, or involves, or consists in, being related to X.
4) By 2) and 3), it looks as though intentionality establishes a relation to the world outside our skins.
5) Intentionality is not a relation.

---

[67] The issue is discussed briefly in (Kriegel, 2011) and (Mendelovici, MS). (Montague, 2013) focuses explicitly on what she call *The Access Problem*, which is quite plausibly taken to be the problem of cognitive contact. However, she seems to assume a relation view of intentionality—though an internalist one.

Clearly, we have a quasi-inconsistent set. I call this set "quasi-inconsistent" because of my use of the 'looks' locution in 2) and 4). This is not intended to leave a theoretical loophole through which a slippery non-relation theorist might slide. My use of 'looks' is merely to indicate that introspection appears to support these theses. I do not intend to make looks-can-be-deceiving arguments.

So, we have a list of theses that, when put together, appear to be inconsistent, since 1), 2) and 3) yield 4), which contradicts 5). Of course, without some clearly defined notion of cognitive contact—a task that I will turn to shortly—1) lends less to the inconsistency of the set than it might otherwise do. However, even with only the murkiest conception of what cognitive contact is supposed to amount to, it is plain that giving up 1) is a consequence to be avoided. 2) Is introspectively obvious: If I am somehow in cognitive contact with, say, the cardinal at the feeder, then it looks as though what is responsible for this contact is my cardinal-at-the-feeder intentional episode. Again, 3) seems almost to follow definitionally from the term 'contact'. I know of no way of contacting anything without being related to it. 4) Is an implication of the acceptance of 2) and 3). As for 5), most would reject it, but, of course, my view is precisely that 5) is true. So what to do?

With some careful analysis, I will suggest that the modificationist is not doomed by the seeming inconsistency. In short, I will suggest that 4) and 5) can be understood as compatible on a proper understanding of 2). What is required here is a clearer picture of 2), such that we can accept that intentionality makes possible our cognitive contact with the world, but is itself not the relation that constitutes cognitive contact. More on this in what follows. For now I want to map out how things will proceed.

First, I want to spell out exactly what an account of cognitive contact should address. Next, I will list some plausible points on which we might evaluate theories of cognitive contact—a set of criteria against which we might judge the relative merits or demerits of various theories. With these in hand, I will go on to examine two views about cognitive contact, evaluating each. I will conclude I will conclude that giving a perfect account of cognitive contact is not an easy thing to do regardless of one's theoretical commitments, and that perhaps all we can hope for respect to cognitive contact is a less than perfect account. I will then explain what options a modificationist might have to provide a plausible account of cognitive contact and explain my preference among these. That being said, I will not claim to have solved the problem of cognitive contact for modificationism. What I hope to do is to show that modificationism is not a doomed theory with respect to delivering cognitive contact, and gesture at how someone might go about working on the problem within the theoretical confines of modificationism.

7.1 What is Cognitive Contact?

Simply put, cognitive contact is the contact that our intentional mental states make with certain things. How, for instance, do my intentional mental states make contact with tables and chairs?[68] To be sure, some theories' account of cognitive contact will just be their account of intentionality. But on other prominent theories of intentionality, intentional mental states are about

---

[68] I will sometimes speak as if it is people who make cognitive contact with things. In a sense, if a particular thought of mine makes cognitive contact with the world, then I too have so contacted it. Strictly speaking, the phenomenon I am after is the contact our intentional mental states make with the world of ordinary objects, so all talk of our/your/my making cognitive contact should be understood as shorthand for intentional mental state talk.

abstract properties, or propositions, so at least some of the cognitive contact we appear to have with concreta stands in need of accounting.

Again, the problem of cognitive contact might seem to bear a family resemblance to the problems discussed in Part 1; namely, the problem of specifying the nature of the relation between intentional states and their contents, and that of specifying what items serve as contents. While the problems are certainly related, they are not identical. The problem of cognitive contact concerns our mental states' relation to certain objects regardless of whether you think that those objects are the contents of thoughts, or instantiate abstract properties that themselves are the contents of thoughts, or, by some other relation, produce mind dependent sense data that serve as contents, or whether you take contents to be modifications/determinants of representational states, etc. We are in cognitive contact with certain objects regardless of your stance on the abovementioned views, and that is what any given theory must account for.[69]

That being said, the problem of cognitive contact and the problem of the nature of content intersect on some views.[70] Direct realists, who take contents to be ordinary objects, will presumably not need to provide an additional story about cognitive contact, apart from their story about how intentional states are related to their contents. On the other hand, our discussion about the abstract objects views of content does not exempt an abstract objects theorist from providing some account of cognitive contact.

---

[69] There are, of course, exceptions. A thoroughgoing metaphysical idealist will have no problem about cognitive contact since his/her ontology does not posit ordinary objects to be in contact with.

[70] That is, the cognitive contact relation just is the intentional relation on some views.

7.2 A theory of cognitive contact

I suppose the most ideal theory of cognitive contact would have the mind somehow extend out in to the world and literally rub itself on the things it appears to be directed at. Assuming the mind is in some way contained in the brain, this does not seem possible.

A close second would be if ordinary objects managed to somehow permeate the skull and, as it were, rub themselves on the mind. As stated, this latter option may also seem problematic. However, I take it that a more refined version is precisely what many theories aim for. Direct realists, for instance, hold that ordinary objects are the constituents of perceptions, and thus that what goes on in perception is partly constituted by ordinary objects and states of affairs. Representationalists, who posit abstract objects as contents, also think the external world manages to imprint itself on the mind.[71] This latter view, however, holds that abstract rather than ordinary objects manage to do the imprinting. This may seem to be an improvement on the ordinary objects view: perhaps something about their abstract nature allows such things to permeate the skull in a way that ordinary objects certainly cannot. However, without some supplemental story about how the imprinting of these abstract objects on our minds manages to secure our cognitive contact, the abstract object view will need to say more about the notion of cognitive contact I am after.

---

[71] Again, talk of imprinting on the mind is talk of something over and above any causal effects the external world has on the mind. With the exception of some ardent idealists, everyone agrees that the external world causes things in us. But this causing is not what constitutes cognitive contact, since the world can cause internal changes in non-cognitive things.

7.2.1 Points of Evaluation

One way to understand the problem of cognitive contact is as a problem about how we manage to *peek* outside our heads and access the world beyond our skins.[72] Understanding the problem this way highlights the closely related issue of what the world beyond our skin is like. That is, being in cognitive contact with the external world would seem to imply some understanding of that world: cognitive contact would seem to obtain when what is going on in the external world and what is going on in the head somehow line up, fit or are otherwise congruent. And that requires a substantial view about what the external world is like. Take, for example, the view that mental representation is a matter of forming a picture in one's mind that resembles the aspects of the external world that are represented. In this case, cognitive contact would obtain when the mental picture is sufficiently similar to the external world (or parts thereof). Clearly, this view assumes a particular view of the external world: It is such that a veridical mental picture can resemble it.

These are deep issues that concern not only theories of mentality, but also metaphysical issues surrounding realism and idealism, and I cannot possibly address everything at issue between realism and idealism on the one hand, and what various theories say about cognitive contact on the other. The point here is to highlight that any account of cognitive contact will, either overtly or by implication, say something about the external world. Hence one way to evaluate theories of cognitive contact is according to how they conceive of the external world with which we are in contact.

---

[72] I say 'peek' but the kind of contact in question is not strictly visual.

Closely related, is what a theory of cognitive contact implies about our knowledge of the external world and the mechanisms that secure it. Ideally, the kind of cognitive contact we have with the external world should, in some way, make possible a relatively robust knowledge thereof. That is, it would be ideal for our theory of cognitive contact to make possible the acquisition of bonafide knowledge about the items we are in contact with in such a way as to leave the mechanism whereby we acquire that knowledge transparent, and clear enough to distinguish bad cases—cases where the failure of this mechanism explains why we fail to acquire knowledge in certain cases.[73]

An ideal account should also either accord with, or give some sort of explanation of, the common-sense view that our cognitive contact with the external world is immediate and direct. At the very least, it certainly *seems* that our cognitive contact with the world is immediate and direct. That being said, one's view need not be that of direct realism. It should, however, explain why direct realism appears to be the unreflective default.

I suppose too that parsimony should be included on our list. My only caveat here is that this last point should figure lower on the list of ideal criteria—a final tiebreaker if you will. My reason for this is simply that I think parsimony is a good explanatory principle, but is not necessarily the gold standard for ontology (Quine and Ockham notwithstanding).

My guess is that there are a whole host of other criteria and ideal cases we might add. For instance, we might want to add that an account of cognitive contact should not appeal to any naturalistically problematic entities or

---

[73] This last point is really a corollary of the bonafide knowledge requirement. Knowledge should ideally be distinguishable from apparent knowledge. I suppose one might be content just knowing that we can have knowledge, even though we might not know we have it.

relations. We would not, for instance, want to count a theory as being ideal if it managed to satisfy all other criteria by endorsing divine occasionalism. A theory that appealed to vital spirits, or the ether, or ghosts, too would be out. I hesitate, however, to include a naturalism proviso, since, assuming that entities that exist outside of space and time, such as Platonic universals, are not natural, this criterion would rule out any view on which such entities played a role in establishing cognitive contact.

7.3 Summary and a Look Ahead

To reiterate then, our evaluation of how various theories understand cognitive contact will be according to the following:
1) How the theory understands the external world
2) What the theory says about our knowledge of the external world
3) How the theory accounts for the seeming directness and immediateness of our cognitive contact
4) A theory's relative parsimony.

At this point, allow me to recap and reiterate how things are going to proceed. The problem for modificationism—the view that intentionality is a kind of phenomenality and that content is a modifying property of phenomenally intentional states—is that it seems unable to deliver, or even make sense of, this intuitive notion that we are in cognitive contact with the world outside our skins. The problem is that it looks like intentionality is what delivers the *relation* of cognitive contact, but modificationism denies that intentionality is a relation. Though a serious problem, I think that the theoretical confines within which modificationism is bound do not spell certain doom for the theory, and that modificationism does have some resources for addressing the problem of cognitive contact. Before

highlighting these resources, however, I want to examine and evaluate how other theories go about addressing the problem of cognitive contact. My methodology for this latter project is, as stated, to see how a given theory stacks up with respect to the points of evaluation listed above. The point, again, is not to enumerate and evaluate every possible theory that has something to say about cognitive contact, but merely to show that the problem is a difficult one. Given this, I will argue that modificationism may indeed have some resources for addressing the problem, though it may turn out that the kind of solution open to the modificationist is, like its rivals, less than perfect.

8: Various Accounts of Cognitive Contact

8.1 Taxonomy

In general questions about our cognitive contact with the world outside our skins are questions about our connections to that world; or, more precisely, about our mental states' connection to that world. How are our intentional states connected to e.g. tables and chairs, such that, for instance, we are able to successfully navigate our environment? How does our cognitive contact with the world differ from how introspection seems to have it?

While it is difficult to draw up lines that neatly categorize various views about our cognitive contact with the world, a good starting place is perhaps to distinguish between what I shall hereafter call *directivist theories* and *indirectivist theories*. Directivist theories take our cognitive contact with the world to be direct and unmediated by entities such as sense data, representations[74], percepts, abstract properties, universals, etc. Indirectivist theories posit something in between our intentional states and the world, such as sense data, representations or abstract objects/properties; and these intermediaries play a role in establishing contact between the mind and the world.

---

[74]My use of 'representations' here might be a bit misleading, since not every one who endorses a representational/intentional theory of mind is an indirectivist. If, however, representations are things that stand in between our minds and the external world of tables and chairs, then a proponent of representations in this sense is an indirectivist.

Later on, it will be made clear where modificationism fits within this taxonomy.[75] Setting this last point aside, a little disambiguation is in order here. First, everyone agrees that intentionality appears to connect us directly to the world of tables and chairs. The directivist/indirectivist divide is thus not about how things seem to an intentional subject. Nor does holding the view that we only see e.g. the table by seeing (some of) its properties necessarily land one in the indirectivist camp because the properties of the table act as intermediaries between one's intentional state and the table.[76] No, the directivist/indirectivist divide categorizes views according to whether or not what is going on in one's mind connects directly to the everyday things it seems to, or whether it connects to something else that is connected to these things via some other relation.

8.2 Indirectivism

As said, indirectivists posit intermediary entities between our intentional mental states and the world of ordinary objects. These intermediaries might be representations, sense data, intentional objects, percepts, or some other kind of thing. Different indirectivist views posit different entities for a variety of reasons. Some, like sense data, are thought to be required because of a disparity between how things appear and how they are. Others are sometimes thought to be required in order to avoid certain complications

---

[75] That is not to say that the details of any given phenomenal intentionalist account will not constrain the possibilities here.

[76] The view in question here is not one that takes the relevant properties to be abstract universals, but one that takes the properties to be property instances. So, one is not an indirectivist simply because one holds the position that we only see the table by seeing **its instantiated** properties. One is an indirectivist if one thinks the table's properties cause us to represent abstract properties, of which the table's properties are instances.

arising from intentional inexistents, and other such puzzles. Below I examine one such indirectivist account, listing its merits and demerits.

Before doing so, I should explain why I have chosen not to discuss a very prominent theory that can be understood as an indirectivist theory of cognitive contact: Frege's account of sense and reference. Roughly put, senses, or cognitive contents on Frege's view (Frege, 1892/1997), act as intermediaries between mental states and the world of ordinary objects. On Frege's view, the cognitive contact relation is akin to reference, which obtains between a mental state and an ordinary object by means of the object's mode of presentation, or sense (Frege, 1892/1997). So for Frege cognitive contact is a two-step relation: the first between the mental state and its content (sense), the second between the sense and its referent (the ordinary object in the present case. My reason for not examining this view of cognitive contact is that I have already discussed several challenges for views that take contents to be senses: if there are challenges to the view that intentionality is a relation to senses, and if this relation to senses is the first of two steps in making cognitive contact (reference), then those challenges apply here as well.

8.2.1 Indirectivism and Sense Data

The sense data theory takes mind dependent concrete objects to be the immediate objects of experience (Russell, 1912/1997, 1927). Though out of favour, sense data theories appear to have some advantages. There does not seem to be a problem about how to specify the relation between intentional states and their contents since the contents with which the intentional states are connected are, at least on most sense data views, in direct contact with them. On the other hand, one might ask how exactly are we to pair sense data

127

with the ordinary objects a theory of cognitive contact requires. Perhaps some sense data theorists are idealists, but many are not. Therefore some additional story about how we cognitively contact the world of ordinary objects by first contacting sense data is required. That is, some account is required about which mind-dependent sense datum is linked to which external ordinary object, such that when I encounter that object, I have that sense datum. That is how a sense data theory would deliver cognitive contact.

One possibility for specifying which sense datum goes with which object is by appeal to co-instantiation. Roughly, round sense data go with objects that have the property of roundness. The challenge here is that sense data theories came about, in part, because our perceptions seem to involve properties that everyday objects do not have, such as perspectival properties. More clearly, sense data theories enjoyed the popularity they did precisely because they had an answer about how perfectly round coins could appear elliptical, how after images could be yellow, etc. The answer was that in such cases, there is indeed something elliptical, or yellow. It is just that the elliptical or yellow thing is a mind dependent sense datum. This is problematic because it seems to preclude the sense data theorist from appealing to the co-instantiation of properties to secure cognitive contact, as the things that we perceive indirectly (the external world objects) often do not, or cannot, instantiate some/all of the properties that the sense data appear to have.

Presumably, someone keen on sense data could posit some causal relation between ordinary objects and minds such that under certain circumstances, encountering object X would cause sense data $_{SD}X$ to happen, token, be instantiated, etc. in the mind of the subject. Whether or not this causal

relation is problematic will depend on the nature of the sense datum. If a sense datum is a physical object, then the causal relation does not seem so mysterious. If, however, sense data are non-physical objects, then the causal relation inherits the problems associated with dualist interactionism: How do physical events cause sense data?

However, because sense data are supposed to have the very properties they appear to, perhaps they are physical objects. After all, assuming sense data theory to be true, and noting that sense data appear to have shape, color, size, etc., they must be physical. Only physical objects have these kinds of properties. But where, then, are sense data located? Russell (1912; 1927), for instance, thought they were located in the brain. The problem with that view is that nothing in the brain is, for instance, blue. On Jackson's (1977) view, sense data are precisely where they seem to be, and are indeed caused by the material objects to which they belong—where *belonging to* is cashed out in the following way:

> ...[A] sense-datum, D, belongs to a material object, M, just if (i) an M-event causes the having of D, and (ii) that spatial properties of D are functionally dependent on those of M as a consequences of the manner in which M causes the having of D. (Jackson, 1977, p. 171)

Jackson also endorses the view that sense data can be three-dimensional (1977, pp. 102-103), which, according to Jackson, provides an answer to Ryle's challenge that "round plates, however steeply tilted, do not usually look elliptical" (Jackson, 1977, p. 103). According to Jackson:

> "[T]he three-dimensionalist has an extra dimension in which to resolve this dilemma. The sense-datum belonging to the round plate held at an angle is round *at an angle.*" (Jackson, 1977, p. 104)

Novel though Jackson's account of sense-data is, one wonders what additional benefit there is to calling the posited entities 'sense data' rather than properties. They belong to objects, are (or can be) three dimensional, have causal influence on our sense organs, and, like the ordinary objects to which they belong, can appear round *at an angle*.

Jackson's idiosyncratic view aside, sense data are often thought to be non-physical—or at least to exist in a phenomenal domain, rather than a physical one—mind-dependent entities that we are directly aware of when we perceive. This being the case, the question that emerges is how sense data theory might deliver cognitive contact. Our points of evaluation were:

1) How the theory understands the external world

2) What the theory says about our knowledge of the external world

3) How the theory accounts for the seeming directness and immediateness of our cognitive contact

4) A theory's relative parsimony.

If sense data exist outside the physical realm, and the ordinary objects that we think populate the external world exist in the physical realm, then it is not entirely clear whether sense data theory has much to say about the external world. Epistemically, sense data theory also faces some serious challenges: our perceptually based beliefs are most plausibly taken to be about ordinary everyday objects and states of affairs. But if perceptions are the justification for such beliefs, and our perceptions are always perceptions of sense data, then we seem to lack proper justification for believing anything about the ordinary objects our beliefs appear to be about—other than, of course, that said objects have, or in some way produce, the relevant sense data. Perhaps this complaint demands too much of sense data theory. The complaint, after all, seems to demand that the theory give us a justification for judgments

130

about the *things in themselves* as it were. On the other hand, the complaint might be taken as merely stating that according to common sense, our perceptually based beliefs are about ordinary objects, not sense data. So, sense data theory owes us an account of why things seem this way. Moreover, to someone not already committed to sense data, it is unclear how positing them would help a theory of cognitive contact. While sense data are concrete particulars on most sense data theories, it is not clear how these mind dependent concreta manage to connect us with the world of ordinary objects. This is especially true if, as is commonly held, sense data are not physical objects, and therefore have properties of a non-physical nature.

8.3 Directivist Theories

I turn now to directivism. Perhaps the most widely known directivist view is direct realism—of which disjunctivism is one prominent species e.g. (Martin 1997, 2004). However, some versions of adverbialism might also qualify as directivist depending on how the adverbial modifications are construed. In any case, what makes a theory directivist is that it posits no entities that mediate between our intentional states and the world of ordinary objects. The relation between our intentional states and the world is, in this sense, direct.

Though direct realism and disjunctivism are theories of perception, I think the issues that arise in the theory of perception are equally applicable to a much broader range of our intentional mental lives. If, for instance, there is a debate about whether your ice-cream perceptions manage to make direct cognitive contact with the ice cream in front of you, then surely the same question could be asked about your desire for said ice cream.

In general, direct realist theories posit a single relation between our intentional states and the world. When I perceive a chair, I do not indirectly perceive it via a sense datum, representation, or set of abstract properties; the chair itself is actually a constituent of my chair-perception. On a very simple direct realist view, the real world, mind-independent chair is a constituent of my chair-thoughts. Cognitive contact is thus automatically established, as the things with which we are in cognitive contact serve as constituents of the very things that are supposed to do the contacting, namely thoughts, desires, perceptions etc. This kind of view appears to accord with how we (unreflectively) think about our cognitive relationship with our environment. Unreflectively, it really does seem as though when I visually perceive the chair, I am in direct contact with it.

8.3.1 Directivist Disjunctivism

In a way, directivist disjunctivist theories—views that claim our veridical thoughts are of a different kind than our non-veridical ones—are a natural progression of direct realism. Because direct realism takes actual wood-and-nail chairs to be constituents of our chair-thoughts, some radically different account of our unicorn-thoughts is required: no flesh-and-blood unicorn was ever a constituent of anything. So, concludes the disjunctivist, hallucinations, illusions, and thoughts about non-existent/impossible objects are not the same kind of thing as veridical thoughts. On such a view, when a subject reports that s/he is currently having a mental episode as of seeing a table, we are to understand this report as saying either that s/he is perceiving a table or that s/he is in some way hallucinating—where the disjuncts name entirely different kinds of mental events/episodes. The report should thus not be taken as indicating that some common mental core is present in both veridical perceptions and hallucinations. For disjunctivists, then, the problem

132

of cognitive contact involves combining direct realism with some plausible story about how subjectively indistinguishable hallucinations differ from veridical thoughts, perceptions, etc. in such a way as to allow that veridical thoughts connect us to the world, but hallucinations, illusions, etc. do not. As a subset of direct realists, disjunctivists claim that the ordinary objects we appear to be in contact with (e.g. tables and chairs) are constituents of our perceptions. The problem is thus to say how this can be so, when a subjectively identical hallucination appears to put us in contact with precisely the same things.

Before evaluating how well disjunctivism fares with respect to delivering cognitive contact, some theoretical preliminaries are in order. As I understand them, disjunctivists do not deny the possibility of subjectively indistinguishable experiences—one veridical, one hallucinatory.

> "The disjunctive account of perception really says that there are two quite different sorts of oasis experience, which may none the less be indistinguishable to their owner" (Dancy, 2009).

> "…it is simply mistaken to suppose that there need be anything more in common across veridical perceptions and delusive experience, other than the fact that all of these states of mind may be indistinguishable for the subject who has them" (Martin, 2009).

What disjunctivists deny is that two subjectively indistinguishable experiences—one veridical, one hallucinatory—are the same kind of mental event. That is, the disjunctivist is not, if the above passages are any indication, denying that a subject may hallucinate a lemon such that, from the subject's point of view, the hallucination is indistinguishable from a veridical

133

case. The disjunctivist claims that, indistinguishable though the two cases seem to be from the subject's point of view, the two cases are not the same type of mental event.

There is something very commonsensical about direct realism, and to my mind, disjunctivism is a very natural progression of the view. If you think that the table in front of you is a constituent of your perception, so that your perception is partly the perception it is because of the real wood-and-nail table, then it follows pretty straightforwardly that a hallucination of a table cannot be the same sort of mental event, since there is no table to make that mental event a perception of a table.

That being said, disjunctivism is often criticised on the grounds that it lacks a satisfactory explanation about what happens in the case of hallucination or illusion. To repeat, disjunctivism is the view that perceiving a lemon, and having a subjectively indistinguishable hallucination as of a lemon, are two distinct kinds of mental events. The former is the perception it is (and a perception at all) because of the presence of an actual lemon. The latter is something else entirely. As Mark Johnson puts it:

> The Disjunctive View has nothing satisfactory to say in answer to the pressing question: What kinds of things can visual experience be a relation to so that in a transition from a case of visual hallucination to a case of seeing there need be no difference which the subject can discern? (Johnson, 2009, p. 216)

While there are numerous arguments in the literature, both for and against disjunctivism, my concern is not so much with the viability of the view as a theory of perception, but as a possible explanation of cognitive contact.

Moreover, what I will have to say is not meant to be devastating to disjunctivism. Rather, what I want to show is that disjunctivism does not have a perfect theory of cognitive contact. My strategy, unlike most critics, is not to press the disjunctivist about what happens in the bad (hallucinatory) cases, but to focus instead on the good (veridical) ones.

8.3.2 A Worry About Subjective Availability

Take some veridical experience such as that of seeing a table. Remaining neutral on what exactly the contents of such an experience are,[77] we may ask what exactly is available to the experiencing subject. Intuitively, what is available to the subject are the table and its features: its shape, color, size etc. Now, imagine a subjectively indistinguishable hallucination. To repeat, disjunctivists do not deny that one might have such a subjectively indistinguishable experience. Again, we may ask what is available to the subject here. Presumably what is available to the subject is precisely the same as in the case of veridically perceiving the table. But, claims the disjunctivist, the two experiences are of a different kind: one involves the table, the other, obviously, does not. Let us assume that the term 'involves' in the previous sentence denotes cognitive contact, so that in the veridical case, cognitive contact with the table obtains. The question that emerges is: What about the good case secures the cognitive contact?

According to Byrne and Logue, disjunctivists hold that "…the good case and the (hallucinatory) bad case share no mental core" (Byrne & Logue, 2009, p. x). This, combined with the comments from Dancy and Martin above, yields the following line of reasoning:

---

[77] Remaining neutral, that is, on the question of the ontology of the contents i.e. abstract objects, sense data, ordinary objects and their properties, etc.

1) One can have a hallucination as of a table that is subjectively indistinguishable from a veridical perception (of a table)—i.e. what is available to the subject is the same in both cases.
2) The hallucination and the veridical perception share no mental core.
3) What is available to the subject is not part of the mental core in either case.

To be sure, a lot will depend on what exactly this notion of a mental core amounts to. However, I think it is fairly natural to understand some mental state's mental core as being something like its identity conditions. If this is right, then disjunctivism seems to divorce what is subjectively available from the identity conditions of the relevant mental states. In other words, one should be able to specify that perceiving the table is the kind of thing it is— namely, a mental state that makes cognitive contact with the table—without mention of what is available to the subject of that state. What is subjectively available is not a factor in establishing cognitive contact with the table.

Why exactly is this problematic? Consider your visual experience of the table in front of you. Does it not seem as though what is subjectively available to you—how the table and its features seem to you—will figure in an account of the cognitive contact you have with the table? Say, for instance, that the table is rectangular and brown, and it also seems to you in a subjectively available way that the table is such. I am not alone in finding it highly implausible that what is subjectively available to you—how the table seems—plays little role in the cognitive contact you have with the table:

One is in causal, sensory, and indeed visual contact with a garden shed, but when one looks at it one has—due to a disorder in one's visual

system, or a distortion of the atmosphere in which light-waves that would have reached one from the shed do reach one, but profoundly rearranged—an experience of a pink elephant. In this case, too, one can locate and track the shed, in spite [of] one's inaccurate conception of it. Does one see the shed? My intuition is that one does not, because one's apprehension…of it is simply too inaccurate. …[C]ontact with an object, in the present sense, is not enough…to guarantee that one sees it, that one is in…contact with it… (Montague, 2013)

For Montague then, cognitive contact with an object requires what is subjectively available to play an important role in the relation. What is subjectively available must have some degree of congruence with an object in order for the subject to be in cognitive contact with it—what is subjectively available does some work in making cognitive contact. And the worry I am raising for the disjunctivist is that it is unclear whether s/he can accommodate this fact.

I want to stress that the worry I am raising here is not just an appeal to internalism. I am not arguing that disjunctivism is wrong because it cannot accommodate certain internalist principles. The worry about subjective availability is consistent with a broadly externalist view. In the present context, an externalist view would be that what is subjectively available is not sufficient to determine what item the subject is in cognitive contact with. What is required by externalism is that the subject's relation to her environment be partly determinative of what, if anything, the subject is in contact with. When Ed, the earthling, and Ted, his twin earth counterpart, are staring at a glass of H_2O and XYZ respectively, what is subjectively available to them (which is indistinguishable) is insufficient to determine that Ed is staring at water, and Ted at twin water. As the story goes, Ed is

staring at water because of the relation he bears to the glass of water in front of him, likewise for Ted and twin water. What externalism claims here is that the identity conditions of Ed and Ted's respective perceptions are not exhausted by what is subjectively available to them, the identity conditions also include the relevant relations they bear to the things in their environment. In short, externalism does not eschew what is subjectively available from the "mental core" of a perception; rather, it says that the mental core consists of more that just what is subjectively available—namely the actual relations the subjects bear to the things in their environments. Disjunctivism on the other hand does seem to exclude what is subjectively available from the mental core.

Let us consider how disjunctivism fares with respect to the points of evaluation listed above. The points on which I suggested we might evaluate a theory of cognitive contact were:

5) How the theory understands the external world
6) What the theory says about our knowledge of the external world
7) How the theory accounts for the seeming directness and immediateness of our cognitive contact
8) A theory's relative parsimony.

As a species of direct realism, disjunctivism would seem to understand the external world as being populated by mind independent objects that are the kinds of things that can be directly perceived. Tables and chairs are members of this population and are presumably, according to direct realism, as we perceive them; though it is unclear whether the disjunctivist can help themselves to this latter claim, given that the notion of 'as we perceive them' seems tied to what is subjectively available.

One might expect a direct realist theory to be among the best contenders for delivering the kind of robust epistemology that one might hope for from a theory of cognitive contact: what better for securing knowledge of the external world of ordinary objects than an account that connects us directly to those objects. The worry here is closely related to the problem of subjective availability. If what is available from the subject's point of view is not involved in establishing cognitive contact, and what is available from the subject's point of view are things such as color, shape, texture, etc., then it is unclear how the latter could deliver the kind of robust knowledge it seems well positioned to deliver.

Setting aside these kinds of concerns, there is little doubt that disjunctivism can easily account for the seeming directness of our cognitive contact with the world: it is not just a seeming; the contact really is direct. Disjunctivism also surpasses the indirectivist in its explanatory parsimony. Without intermediary entities in the cognitive contact relation, the disjunctivist need not provide some extra account of how the contents of our intentional states map on to the world of ordinary objects: ordinary objects just are the contents. On the other hand, without some plausible account of what happens in the bad cases, that makes intelligible the claim that such cases are of a different kind than the good cases, this parsimony may not confer much of an advantage.

8.4 Tallying Up the Points

What I have tried to do in this chapter is to show that cognitive contact is a difficult problem for both directivists and indirectivists alike. Neither of the views we considered had a perfect theory of how we manage to make

139

cognitive contact with the world of ordinary objects, and I am inclined to think that this is due not so much to the shortcomings of the respective theories, but to the recalcitrance of the problem. I chose to look at the two views I did because they serve as exemplars of the two approaches to addressing the problem—directivism and indirectivism. In general, directivism appears to best accord with the common sense view, but disjunctivism in particular, appears not to afford the mind much of a role to play in establishing cognitive contact. Sense data theory, on the other hand, appears to have the opposite problem: the mind and its objects, sense data, appear cut off from the world of ordinary objects. All in all, the point of this survey was not to refute any particular theory, but rather to show that a perfect account of how our minds make contact with the world is perhaps too much to hope for. We should thus accept that what we can have with respect to cognitive contact is a less than ideal account. Given more modest aspirations for a theory of cognitive contact, I think the options I will provide below will be seen to fare as well as their rivals.

9 Modificationism and the Challenge of Cognitive Contact

As mentioned above, modificationism is the combination of the moderate identity view—intentionality is a kind of phenomenality—with the view that phenomenal character is a way for a mental state to be—a modification of that state. While modificationism appears not to make a claim about cognitive contact, some analysis ought to at least constrain the view's possibilities. It is precisely these possibilities that will be the focus of this final chapter. To reiterate, we began with a set of sentences that appear inconsistent:

1) We are in cognitive contact with the world outside our skins.
2) It looks as though this contact is established by intentionality.
3) Being in cognitive contact with X is, or is partly constituted, or involves, or consists in, being related to X.
4) By 2) and 3), it looks as though intentionality establishes a relation to the world outside our skins.
5) Intentionality is not a relation

As I mentioned earlier, the task for the modificationist is to explain how a non-relational phenomenon such as intentionality can in some way establish a relation like cognitive contact. In what follows, I would like to sketch several possible ways a modificationist might address this challenge.

9.1 Two Kinds of Intentionality:

One avenue a modificationist might explore in the pursuit of cognitive contact is to endorse two kinds of intentionality. In (Horgan, Tienson, & Graham, 2004), a view is sketched according to which there are two kinds of

intentionality: phenomenal and externalistic. As is the case with modificationism, "[p]henomenal intentionality is narrow: it is not constitutively dependent upon anything outside the head of the experiencing subject. Indeed it is not constitutively dependent on anything outside of phenomenal consciousness itself" (Horgan, Tienson, & Graham, 2004, p. 299). On the other hand, there is a kind of intentionality, externalistic intentionality, which does depend on things in the phenomenal subject's environment.

For the modificationist, this might seem a rather big concession: modificationism is a view according to which intentionality is a kind of phenomenality. Since phenomenality is not constitutively dependent on anything outside the experiencing subject (i.e. it is narrow), allowing a kind of externalistic intentionality seems to give up the thesis that intentionality is a kind of phenomenality.

However, I think the modificationist has a plausible reply here. First, this notion of externalistic intentionality might be construed as something like reference. The modificationist might therefore argue that there is content on the one hand, and reference on the other, and that phenomenal intentionality concerns the former, but that cognitive contact concerns both. The idea, according to Horgan et al, is as follows:

> "Suppose that you have an occurrent thought that you could express linguistically by "That picture is hanging crooked," where the singular thought-constituent expressible linguistically by 'that picture' purports to refer to a picture on the wall directly in front of you. This thought-content involves certain phenomenally constituted presuppositions, which we call *grounding* presuppositions, that must be satisfied in order for…the

142

thought to refer: roughly, there must be an object at a certain location relative to oneself …, this object must be a picture, and there must not be any other picture in that location that is an equally eligible potential referent of 'that picture'. If these grounding presuppositions are satisfied by some specific concrete particular, then your singular thought constituent refers to that very object. Which object your thought-constituent refers to, if any, thus depends jointly upon two factors, one phenomenally constituted and one externalistic: on one hand, the phenomenally constituted grounding presuppositions, and on the other hand, the unique actual object in your ambient environment that *satisfies* those presuppositions. " (Horgan, Tienson, & Graham, 2004, p. 305)

In terms more consonant with modificationism: The grounding presuppositions are part of the phenomenal intentional content of one's crooked picture thought, and one's crooked picture thought makes cognitive contact with the picture when the picture satisfies this content. Of course something more is needed here, since my twin earth phenomenal duplicate will, ipso facto, enjoy type-identical phenomenally constituted presuppositions. And since there is a twin picture on this twin earth, what is it about my and my twin's respective thoughts that ensures that I am in cognitive contact with *this* picture, and he with *that* picture? This is where allowing some externalistic elements to figure in the relation of cognitive contact is helpful. It is facts about the actual picture in my environment and my relation to it that ensure my cognitive contact with *this* picture, likewise for my twin.

According to Horgan et al, many thoughts have two kinds of truth conditions, wide and narrow. The narrow truth conditions for the thought 'that picture is crooked' are those you share with your twin earth and brain-in-a-vat

counterparts: "there is a unique object x, located directly in front of me and visible by me, such that x is a picture and x is hanging crooked (relative to my visual/kinesthetic up/down axis)" (Horgan, Tienson, & Graham, 2004, p. 313). The wide truth conditions for this same thought have as constituents, the actual satisfiers of the thought. Since your crooked-picture thought and your twin earth duplicates have different actual satisfiers, your respective thoughts have different wide truth conditions.

Recall that our inconsistent set of sentences were:
1) We are in cognitive contact with the world outside our skins.
2) It looks as though this contact is established by intentionality.
3) Being in cognitive contact with X is, or is partly constituted, or involves, or consists in, being related to X.
4) By 2) and 3), it looks as though intentionality establishes a relation to the world outside our skins.
5) Intentionality is not a relation.

The kind of two-factor approach advocated by Horgan et al might be co-opted by the modificationist as a means of reconciling the seeming inconsistency. The modificationist intent on this kind of account would presumably argue that cognitive contact is not solely established by phenomenal intentionality (as per 2), but by phenomenal intentionality + externalistic factors (call this 2a). Phenomenal intentionality sets conditions that externalistically determined referents then satisfy. Or to put things another way, cognitive contact is the relation that obtains when both narrow and wide truth conditions are satisfied. In short, cognitive contact is established partly by narrow phenomenal intentionality—its narrow phenomenally constituted presuppositions and truth conditions—and partly by the wide truth conditions. Modificationism is a story about the former, which, when

combined with a suitable account of wide truth conditions and reference, can provide a plausible account of cognitive contact.

9.2 Descriptive Space, and Phenomenal Flitting

As a reminder, modificationism holds that my red-table intentional episode consists in my being intentionally directed in a red, table-esque way. This kind of view bears some similarity to the adverbial theory of perception, as advocated by theorists such as Chisholm (1957). Adverbialists held that intentional experiences are not, as the sense data view claims, relations to contents, but instead modifications of the intentional experiencing subject. Where the sense data theorist invokes red, round sense data to explain our tomato experiences, the adverbialist contends that we are experiencing redly, and roundly. Rather than being independently exiting things (properties, sense data, etc.), the adverbialist takes redness and roundness to name the way a subject is intentionally experiencing. Adverbialism does not make a direct claim about how experiencing redly might make cognitive contact with the tomato, but there are options.

One avenue a modificationist might take is to provide some story about how representing redly, roundly, etc. might, in some cases, determine a kind of descriptive space sufficient to establish cognitive contact. On such a view, when I represent redly and roundly, my intentional state sets conditions—i.e. determines a descriptive space—that may or may not be satisfied/fit by objects in my environment. Here too, modificationism shares some similarity to another theory from the recent history of philosophy, namely *descriptivism*. And given that there are well known challenges to descriptivist views of content, if the modificationist is to avail herself of descriptivism in her account of cognitive contact, these will have to be addressed.

First however, more will need to be said about the kind of determination invoked in the phrase 'determine a kind of descriptive space'. Moreover, the notion of descriptive space—especially to someone who has been at pains to argue that intentional contents are phenomenal modifications—stands in need of explanation: e.g. What exactly is this descriptive space; is it a kind of content; what is its relation to intentional content; is it relational, propositional, similar to reference, sense, extension, intension etc.; is it, as per indirectivism, something in between our intentional states and ordinary objects?

To begin with, I want to try to clarify how intentional content can "determine" a descriptive space. Several theorists, (Kriegel, 2011. Mendelovici, MS. Searle, 1983), have noted that just from how things are for me when I undergo some intentional experience, I have a pretty good idea about how the world would have to be in order for my intentional experience to be accurate, true, satisfied, etc. Of course, I do not mean to imply that any given intentional experience is sufficiently detailed so as to specify the state of the entire world at any given moment. But then again, intentional experiences do not typically purport to be about the entire world. When I think to myself that this was a long winter, when I see my azaleas in full bloom, when I desire another rescue dog, I have a pretty good idea how the world would have to be in order for these experiences to be satisfied: the world would have to be just like how my thoughts have it. Moreover, it is not immediately obvious that I need look any further than how things are for me from the skin in to determine how the world would have to be. If, for instance, Descartes' evil demon were to somehow take a mental snapshot of what is going on inside me, he would know precisely how to arrange the world for me to be veridically encountering it.

If we call how the world would have to be in order for my intentional episodes to be satisfied the descriptive space determined by an intentional experience, then this notion of intentional episodes determining a descriptive space becomes clearer. How things are for me, intentionally, determines how the world would have to be in order for my intentional states to be satisfied. Or, to put things in a way more consonant with the spatial metaphor, how things are for me determines a descriptive space into which the world may more or less fit.

Again, this approach bears some similarity to the descriptivist theories championed by authors such as Frege, Russell and later Searle. As developed in (Frege, 1892/1997), (Russell, 1905/2008) and (Searle, 1958), descriptivism is a theory about proper names and how they function: Descriptivists hold that in addition to their reference, proper names also have a meaning, sense or descriptive content, and that it is in virtue of this that they refer. Searle, for instance, contends that a name refers by being associated with a cluster of descriptions, a vague, unspecified number of which are true of the thing referred to (Searle, 1958).

Though modificationism is a theory about the nature of intentionality and phenomenality, not language, the current solution to the problem of cognitive contact we are examining on behalf of the modificationist is one that avails itself of the notion of a descriptive space. What I have in mind by 'descriptive space' is similar to what I take Searle to have in mind with his cluster theory. For Searle, a proper name refers by being associated with a cluster of descriptions that the referent satisfies; for the modificationist, intentional episodes make cognitive contact by determining a descriptive space into which the world (or objects therein) more or less fit.

But how exactly does this descriptive space arise? How does content determine a descriptive space? Consider the following example: The last time I spoke with my sister, our conversation went from the usual greetings, to my thesis, to my brother-in-law's most recent fishing exploits, to the puppy my wife and I recently got, to my sister's kids, etc. As is always the case when speaking with my sister—whose many qualities include ranking among the world's best moms—she asked certain questions that lead me to believe she was fishing for a sense of how my wife and I were coping with the stress of writing, and for any clue as to how she might help alleviate some of it. Besides appreciating my sister's concern, the sense that she was fishing for such insights led me to ask about any recent fishing adventures her husband might have had. At sometime during my realization that my sister was fishing for insight into my stress level, there was a brief, phenomenal presence of my sister's husband—owing no doubt to my thinking that my sister was fishing for insight—that flitted across my mind and led me to ask about her husband's fishing. When my sister began asking about our new puppy, there were several moments that the thought of her kids flitted through my mind. In particular, for whatever reason, I remember clearly the look of awestruck wonder on her son Felix's face when he was presented his first toy wheelbarrow so he could help his uncle garden. Perhaps this is why, when I asked about her kids, I asked whether Felix has been helping my Dad do the gardening. Again, I assume the most plausible explanation here is that my sister was asking about the closest thing my wife and I have to kids, our dogs, and that led me to think about her kids, Felix and Charlotte. What I am trying to convey is that such phenomenal flittings contribute in an important way to the determination of the descriptive space.

Another feature of the descriptive space is its fluidity: It is not some rigid set of descriptions, finite and exhaustively expressible in language. It is under

frequent modification by the phenomenal flittings that serve to determine it. Horgan, Tienson and Graham make similar remarks about narrow content and the grounding presuppositions that figure therein:

> "…[N]arrow truth conditions…are not compactly formulable linguistically in a way that can neatly be plugged into the right side of statements of the form "Statement 'S' is true iff…" [T]he background presuppositions figuring in the narrow content of intentional mental states typically cannot be spelled out in any tractable way, and…these presuppositions typically are too complex and ramified to be cognitively surveyable…" (Horgan, Tienson, & Graham, 2004, p. 314).

Several questions/concerns arise here: First, how well do objects need to fit the descriptive space in order for cognitive contact to obtain? (This concern has an analogue in the descriptivist/referentialist debate in the philosophy of language: how many of the descriptions must something satisfy to be the referent?) Second, though originally about proper names, Kripkean-type objections can be adapted to apply here as well. One of Kripke's original insights was that someone could refer to Gödel by 'Gödel', even though none of the referrer's descriptions of Gödel are true of him (Kripke, 1972). Hence description cannot be the mechanism that secures reference.

The same might be said of the modificationist's descriptive space. There might be cases where, a) though some object does not fit the descriptive space determined by a subject's phenomenal intentional content, we are nevertheless inclined to attribute cognitive contact between the subject and object, and b) where some object fits the descriptive space, but is not something the subject is in cognitive contact with. Examples of a) include imposter cases. Imagine one's wife is replaced with a cleverly disguised

robot. Though the imposter robot will fit some of the descriptive space—the part that involves superficial surface features (e.g. hair color)—it will not fit other parts. The phenomenal flittings that often occur when interacting with one's spouse—including flittings of emotion, past events, etc.—will undoubtedly contribute to the determination of a descriptive space into which the imposter will not fit. The point is that, were one to have lunch with the imposter robot, there is a pretty clear sense in which one is in cognitive contact with the imposter despite it not fitting the descriptive space. Examples of b) might include cases of type identical, twin earth artefacts: It could be argued that both the pencil in your hand, and its twin earth counterpart, fit the descriptive space that secures the cognitive contact between you and the pencil. But surely you are not in cognitive contact with the twin-pencil.

Let us first address a): How might a modificationist address the imposter objection? To repeat, the objection is that there are reasons to think that when having lunch with a robot that is cleverly disguised as one's wife, one is in cognitive contact with the robot. This runs counter to the modificationist's claim that the mechanism by which cognitive contact obtains is fit: objects fit the descriptive space determined by phenomenal intentional content (including phenomenal flittings). Since the descriptive space determined in this instance would include flittings of emotion, past events, etc., and since the robot does not satisfy these flittings (your emotions are not robot-directed, your wedding did not involve the robot), the robot does not fit the descriptive space. But, the objection continues, one is in cognitive contact with the robot. Ergo, cognitive contact is not, or not just, a matter of descriptive space.

In response, the modificationist might simply deny that one is in cognitive contact with the robot. The robot does not fit the descriptive space,

therefore, incredible though it undoubtedly seems, one is not in cognitive contact with the robot with which one is having lunch. Less incredibly, the modificationist might contend that the contact one has with the imposter is degenerate, or somehow divided. The idea here would be that the descriptive space determined during your lunch with the imposter is partially satisfied by the robot, but also partially satisfied by your actual wife. That being the case, there is something amiss, cognitively speaking, about the contact one has with the robot. The robot fits certain superficial surface features involved in the descriptive space, but fails to fit other elements such as those determined by flittings of emotion and past events. In my view this is just the kind of response one would want in this situation: the contact you have to the thing you are having lunch with is imperfect: You take yourself to be having lunch with your wife, but are instead having lunch with something that satisfies some descriptions of your wife, but not others.

With respect to b), the modificationist also has a couple of possible responses. To repeat, the objection was that a pencil on twin earth would equally fit the descriptive space determined during my interaction with the pencil on my desk. But I am in cognitive contact with the pencil on my desk, not the type identical pencil that exists in a different world. Therefore, the fitting of a descriptive space cannot be the mechanism that establishes cognitive contact. Again, the modificationist might simply argue that, implausible though it seems, cognitive contact can indeed be a trans-world relation. Notice too that the modificationist cannot avail herself of the kind of emotional, or past experiential, flittings that are involved in the cognitive contact one has with one's spouse. Odd fetishes notwithstanding, we do not form those kinds of bonds with pencils. An interesting avenue that a modificationist might pursue here is the token reflexive indexical account proffered by Searle's theory of intentionality (Searle, 1983, pp. 218-225). In

some cases, phenomenal intentional content determines a descriptive space that includes a token reflexive indexical component. In our pencil case, this component can be glossed as 'the cause of *this* very phenomenal intentional episode'. When undergoing a phenomenal intentional episode as of the pencil, the descriptive space determined, includes the component 'the cause of this very phenomenal intentional episode'. Since the twin pencil does not fit that part of the descriptive space, one is not in cognitive contact with the twin pencil.

But there is a lurking worry here. The modificationist is committed to intentionality being a kind of phenomenality, and being non-relational. Phenomenal intentional content is a way for an intentional state to be. The solution (that we are now considering) to the problem that such a position engenders—namely, the problem of cognitive contact with the external world—is that phenomenal intentional content (including phenomenal flittings) determines a descriptive space into which objects in the external world more or less fit. This led us to the present objection that some twin-world artefacts will fit the descriptive space determined by some phenomenal intentional contents. So by the theory's own lights, we are in cognitive contact with twin-world artefacts. The current proposal for answering this objection is to invoke, as part of the descriptive space, a self-reflexive indexical component that, in the present example, is something like 'is the *cause* of this very phenomenal intentional episode'. But how can a theory that has been at pains to eschew relations such as causation from being constitutive of intentionality (see Part 1) appeal to that very relation to get out of twin-world hot water (or XYZ as it were)? In other words, isn't the modificationist helping herself to something she cannot have?

To answer this concern on behalf of the modificationist, let me start by saying that I do not think any modificationist denies that the world outside our skins can, and regularly does, have a causal impact on us, and that often the effect—by complicated psycho-physical processes—of this impact is the production of phenomenal intentional content. What the modificationist denies is that this content is constitutively dependent on the causal relation. It is contingently true that pencils often cause pencil-esque phenomenal intentional episodes, but such an episode could occur in a brain-in-a-vat, and the modificationist is committed to the view that the latter is every bit as pencil-esque as the former. What makes it the kind of phenomenal intentional episode it is is not that it is pencil-caused (though it may have been), but that it has certain phenomenal features that we can gloss as 'pencil-esque' (again, what this actually comes to for the modificationist was addressed in part 2). Now, consider how the modificationist would account for having a phenomenal intentional episode as of kicking. Presumably, she would say that what it is to have a kicking-thought is to be intentionally directed in a kicking-wise way. Does this, much to the modificationist's chagrin, commit the modificationist to a relational view? In tamer language: Kicking is a relation. A subject can represent kicking. Does that mean the representation is relational?

To come back to b)—the objection that twin-world objects fit worldly descriptive spaces—I do not think it is straightforwardly inconsistent for the modificationist to contend that, in the case of the pencil, the reason why the twin-earthly pencil does not fit the descriptive space is that part of the descriptive space determined includes a token-reflexive component. This kind of view might also help the modificationist account for the following kind of brain-in-a-vat (henceforth, BIV) scenario: Your BIV duplicate's phenomenal intentional contents will determine a descriptive space into

which no pencil fits. This is because his descriptive space will include a token reflexive component that may be glossed as 'the pencil causing **this** very pencil-esque intentional episode', and nothing fits that space. I find this kind of view promising, but my guess is trying to spell out the phenomenal intentional content in a coherent way will be challenging. Again, the situation is this: Your current phenomenal intentional episode as of the pencil needs to be explicated such that it is clear how it might determine a descriptive space that guarantees your cognitive contact with the pencil, and not its twin-earth counterpart, and also makes clear how your BIV duplicate's similar phenomenal intentional episode does not make cognitive contact with either object. The modificationist story here is that your current pencil episode has certain phenomeno-intentional properties that may be glossed as: yellow-pencil-causing-this-very-experience-esque. That is, you are intentionally directed in a yellow-pencil-causing-this-very-experience-esque sort of way. On the one hand, these kinds of translations always come across as clumsy and contrived. On the other hand, it really does seem like many phenomenal intentional episodes—especially perceptual ones—are token-reflexive in this way.

Coming back to the question of how well something needs to fit the descriptive space in order for cognitive contact to obtain, I want again to draw an analogy from the philosophy of language. According to Searle, the criteria for applying a proper name are loose, and unspecified, rather than rigid (Searle, 1958, p. 172). By this, I take him to mean that the descriptions some object must satisfy in order to be the referent of a proper name is a loose set with an unspecified number. Likewise for the descriptive space and the object that fits it. In her account of how exactly we access (perceive/think about) external objects, Michelle Montague expresses a similar view:

"We achieve access to a material object via perception by correctly representing enough of that object's properties. It is difficult if not impossible to give a principle for determining when there is enough matching for perceptual contact" (Montague, 2013, p. 46).

Again, the problem of specifying how well an object must fit the descriptive space may be partly due to the fluid nature of the descriptive space itself. On the modificationist line we are examining, the descriptive space is being frequently modified, adjusted, and revised. The thing about phenomenal flittings is that they flit—making their contribution to the descriptive space before giving way to other flittings, which subsequently modify the descriptive space too. When I think of my Pyrenees, Sibyll, on one occasion, and then on another, the chances are that some of the phenomenal flittings that occur on each occasion will be different. Sometimes the thought of how parental she is with our cat flits through my mind. Other times, it is the maddening frequency with which she insists on rolling in mud (she's our only white dog).

Before moving on, a brief summary of this descriptive space view and how it addresses the problem of cognitive contact will be helpful. According to the present line, phenomenal intentional content, which includes phenomenal flittings, determines a descriptive space—a way-the-world-would-have-to-be. Cognitive contact is achieved when ordinary objects fit, more or less, into the space. How well an object must fit the descriptive space—i.e. how much of the descriptive space must be satisfied by the object—is difficult to say. This may be due, in part, to the fluid and changing nature of the descriptive space.

Recall our inconsistent set of sentences:

1) We are in cognitive contact with the world outside our skins.
2) It looks as though this contact is established by intentionality.
3) Being in cognitive contact with X is, or is partly constituted, or involves, or consists in, being related to X.
4) By 2) and 3), it looks as though intentionality establishes a relation to the world outside our skins.
5) Intentionality is not a relation

The descriptive space view attempts to address the inconsistency by accounting for intentionality in such a way that intentionality is responsible for establishing cognitive contact, without thereby being a relation itself. Phenomenal intentional content determines a descriptive space into which items in the external world fit (more or less). In this sense, the present line understands 2) as saying that phenomenal intentionality establishes the conditions that make cognitive contact possible (assuming the world cooperates) by determining this descriptive space.

9.3 Descriptive Space, Directivism and the Externalism Issue

Recall that one of the points by which we evaluated theories was how the theories account for the seeming directness and immediateness of our cognitive contact with the world. This raises the question of whether the descriptive space view is directivist or indirectivist. In other words, does the descriptive space stand in between phenomenally intentional mental phenomena and the world of ordinary objects, and in so doing, make the view indirectivist? Though I take it that advocates of something like the

descriptive space view would opt for directivism, it is worth examining whether the descriptive space view and directivism are really consistent.

Let us start with an analogy. A grocery list, though lacking original intentionality, seems to determine something like a descriptive space, that items in the grocery store more or less fit. The question is whether the descriptive space determined by the list somehow stands in between the list and the items that would satisfy it, such that a theory of grocery contact that assumes this kind of fit mechanism is necessarily indirectivist. To my mind, there is nothing problematic about endorsing this descriptive space view about grocery lists while maintaining directivism about grocery contact; likewise for the descriptive space view of phenomenal intentionality.

That being said, the grocery list analogy can serve to highlight what is undoubtedly the descriptive space view's weakest point. While we have no trouble understanding how a grocery list can *fully* determine a descriptive space, nor precisely what would fit that space, the same is not true of phenomenal intentional content. Though the descriptive space is fluid, and under constant modification, and several theorists recognize that the answer as to how well something needs to fit the descriptive space in order for cognitive contact to obtain is going to have to be vague and unspecified, and probably not something exhaustively expressible in language, this kind of view is unlikely to win any converts. As a phenomenal intentionalist cousin of descriptivism, the descriptive space view's chief opponents will likewise be related to descriptivism's chief opponents—namely externalists of one stripe or another. And just as the externalist who opposes descriptivism presses the descriptivist to provide some cogent account of how many descriptions a potential referent must satisfy in order for reference to obtain, so too will the externalist who opposes modificationism press the modificationist for an answer about how well something must fit the

descriptive space in order for cognitive contact to obtain. Consider Kriegel's scant treatment of the issue:

"One way to think of this is that intentionality provides truth *conditions,* or accuracy *conditions,* which then may or may not be satisfied, depending on the world's cooperation. When the conditions are satisfied, cognitive contact with the world will have been established. The role of intentionality is only to make such contact possible by *laying the conditions* whose satisfaction would constitute the establishment of contact. ...I conclude that when the connection-to-the-world requirement is properly understood, there is reason to expect adverbialism to meet it." (Kriegel, 2011, p. 166)

Assuming the point of view of an objector, we might ask what exactly counts as satisfaction here. That is, how many of the abovementioned conditions must be satisfied for cognitive contact to obtain? Though I am sympathetic towards this kind of descriptive space view, I recognize that solving the problem of cognitive contact is not as simple and straightforward as the cursory treatment of the issue by the phenomenal intentionalist community would suggest: a great deal more work is required here. Moreover, given the shortcomings of many past attempts to ground our contact with the world in sensory terms alone, it is clear that phenomenal intentionality theory in general, and modificationism in particular, ought to take the problem of cognitive contact, and the resources available to address it, seriously. What I have tried to show is that modificationism is not necessarily doomed with respect to cognitive contact, but again, more work is required.

9.4 Taking Stock

So far, we have considered two possible ways a modificationist might go about addressing the problem of cognitive contact. The first was that there are two kinds of phenomena involved in the cognitive contact relation: phenomenal intentionality, and wide intentionality. The second was the descriptive space view. Though not solutions to the problem of cognitive contact, each can be seen as a starting place for addressing the problem. Let us now briefly examine how these views stack up with respect to our points of evaluation, i.e.:

1) How the theory understands the external world
2) What the theory says about our knowledge of the external world
3) How the theory accounts for the seeming directness and immediateness of our cognitive contact
4) A theory's relative parsimony.

With respect to 1), the two-factor theory proffered by Horgan et al appears to be realist in the sense that it understands the external world as populated by the kinds of things that can satisfy the truth conditions of intentional states. True thoughts about crooked pictures will make cognitive contact with actual crooked pictures. Crooked pictures therefore number among the things that populate the external world. Likewise, the descriptive space view also understands in the same way.

Unfortunately, the descriptive space view cannot deliver a robust epistemology: nothing about the descriptive space guarantees that the ordinary objects we are in fact in cognitive contact with actually do fit the descriptive space, nor that we can tell when they do or do not so fit. On the other hand, because the two-factor theory allows some externalistic elements to play a role in cognitive contact, it would seem to have an advantage with

159

respect to condition 2. Externalism is often admired (or criticised) for its anti-sceptical (with respect to knowledge of the external world) implications (see, for instance, (Greco, 2004). Briefly: sceptical arguments often rely on calling into doubt what is introspectively available to the subject. In other words, these arguments assume that justification (and therefore the knowledge it delivers) is an internal matter—concerned only with what is introspectively accessible. The next step in such arguments is to impugn what is introspectively accessible, and thereby threaten justification. However, if it turns out that the justification for some belief depends, in part, on matters external to what is introspectively accessible to the subject, then sceptical arguments that seek to impugn what is introspectively accessible do not necessarily threaten justification, and therefore knowledge. Hence a theory that allows certain external factors to play a role in establishing cognitive contact with the world might be better positioned to give an account of knowledge of that world. To be sure, this scant treatment of the epistemic implications of the two-factor theory is underdeveloped, and blurs the distinction between epistemic and semantic externalism/internalism. However, at the very least, such considerations confer a prima facie epistemic advantage to the two-factor theory. Both the two-factor theory and the descriptive space view are directivist, and so accord with the view that our contact with the world is direct and unmediated by things such as sense data. Finally, with respect to parsimony, both theories appear equal.

To summarize, modificationism has at least two possible avenues for pursuing a solution to the problem of cognitive contact. The first is more theoretically inclusive because of its invocation of some externalist principles. While some phenomenal intentionalists might be content with this first option, the second option—the descriptive space view—is a live option for those phenomenal intentionalists in general, and modificationists

in particular, determined to provide a non-relational, internalist theory of intentionality—one that is, as Katalin Farkas says, without compromise (Farkas, 2008). Again, neither option is meant as a decisive solution to the problem of cognitive contact, but are instead sketches of how modificationists might go about addressing the problem—sketches that I think are within the theoretical constraints within which modificationists, by their own lights, must work.

9.5 Two More *Possibilities*

Before concluding I would like to briefly mention two more approaches to the problem of cognitive contact that a modificationist might take. The first is not so much a possible solution to the problem of cognitive contact as a flat denial that it is a theory of content's job to provide such a solution. The idea here is that phenomenal intentionality theory and its subspecies, modificationism, are theories concerned with psychological content, not semantic relations such as cognitive contact, reference or truth. In the words of Jerry Fodor:

"Truth, reference and the rest of the semantic notions aren't psychological categories. What they are is: they're modes of Dasein. I don't know what Dasein is, but I'm sure there's lots of it around, and I'm sure that you and I and Cincinnati have all got it. What more do you want?" (Fodor, 1981)

Again, I will not belabour this point too much because my focus is on the possibilities the modificationist has for facing the problem of cognitive contact, not how she might avoid it.

The final avenue down which a modificationist might venture in pursuit of some way to address the problem of cognitive contact is what I dub reconstructive realism. Before going briefly into the view, a caveat: The following is but the barest of sketches—one that I have spent very little time on, but one that strikes me as extremely interesting. It was suggested to me that what it is to be an object for us is to be something like the locus of engagement, or interaction possibilities[78]. This view has some intuitive force: the reason why the thing in my left hand is a pen is because of certain ways I can engage or interact with it. And for this very reason, the thing in my left hand cannot be a pen for my dog. It can be a fetch toy, a stick (and most likely a chew toy to be shredded into inky pieces on my new duvet). But it cannot be a pen for my dogs. Of course, certain logical snafus are bound to arise. For instance, if I throw the pen for my dog, then what I throw and what he fetches are different objects. Perhaps this kind of problem can be resolved by appeal to overlapping possibilities. The pen can also be such that I can engage it as a fetch toy for my dog, and so there is some overlap in how I can engage the pen, and how my dog can. At any rate, we are bidden, according to this line of thought, to reconstrue the world of ordinary objects as loci of interaction/engagement possibilities (Bickhard, 2010).

Undoubtedly, this kind of view smacks of idealism: Ordinary objects are not defined independently of our impressions of them, but are defined instead in terms of how we might interact with them. Maybe this sounds too far-fetched to be a metaphysical account of ordinary objects, but consider the phenomena of invention and discovery. Ancient man needed some way of transporting heavy things over great distances. One day, someone noticed that it was easier to move fallen trees by rolling them rather than lifting them, and all of a sudden circular objects became more than just sections of

---

[78] By Chris Viger in conversation

fallen trees...they became wheels. What were loci of engagement possibilities that included, for instance, fuel for the fire, poles for simple hut construction, etc. became new things; they became wheels.

Though this is a mere sketch, it might be further fleshed out along similar lines to the enactive approach to perception developed in (Noë, 2004). Briefly, Noë's view is that perception is something we do, not something that happens to us. "Think of a blind person tap-tapping his or her way around a cluttered space, perceiving the space by touch, not all at once, but through time, by skilful probing and movement" (Noë, 2004 p. 1). The world, according to Noë, "makes itself available to the perceiver through physical movement and interaction" (Noë, 2004 p. 1). To see the tree over there is to see something up which one might climb. "It is to see it, directly, as affording certain possibilites" (Noë, 2004 p. 106).

To be sure, modificationism and Noë's enactive model are incompatible in many other ways. For instance, according to Noë, one implication of the enactive approach is that we ought to "reject the idea—widespread in both philosophy and science—that perception is a process in the brain whereby the perceptual system constructs an internal representation of the world" (Noë, 2004 p.2). But that need not prevent the modificationist from adopting certain Noë-esque metaphysical views about the nature of ordinary objects. The important point is that there is this view of what it is to be an object: To be an object is to be the loci of interaction possibilities, or, in Noë's words, to be such as to afford certain possibilities. And this view might be co-opted by the modificationist in her account of cognitive contact.

How exactly could reconstructive realism help the modificationist in her pursuit of some account of cognitive contact? Well, all along, the assumption

has been that the problem of cognitive contact is a problem about how to account for what's going on in the mind of a cognizer such that what is going on there manages to reach out into the world and make contact with things. But on the present view, our focus ought to be on how we conceive of, and account for, the things we take ourselves to be in contact with, such that they can be possible candidates for cognitive contact. In slightly more earthy language: The present line contends that the problem of cognitive contact is not a problem about how we fix the mind such that it can contact the world, but about how we conceive of the world such that it is the kind of thing that could be in contact with the mind. And indeed, how you might go about engaging/interacting with something seems, at least prima facie, to be the kind of thing for which one's phenomenal intentional experiences could play a central role. Perhaps reconstructive realism could even be combined with the descriptive space view such that what it is to determine a descriptive space is to determine a space of interaction possibilities. Again, I realize this is extremely underdeveloped and exceedingly vague, but for the modificationist, it might be worth pursuing.

9.6 Conclusion

The central goal of this project was to examine what I find to be an interesting family of views that fall under the common head of phenomenal intentionality theory. More specifically, I wanted to examine the strengths and weaknesses of a particular species of phenomenal intentionality theory advocated by theorists such as Kriegel (2011) and Mendelovici (2010). This species of theory, which I called modificationism, rejects the view that intentionality is a relation, but is instead one kind of another pervasive mental phenomenon, phenomenality.

Before examining the strengths and weaknesses of modificationism, I examined what motivations might lead someone to reject the relational view of intentionality. I looked at what are undoubtedly the most popular answers to two questions:

1) What sorts of things does intentionality relate us to?
2) What kind of relation is intentionality?

Though certainly not an exhaustive vetting of all relational views, I think the concerns I raised—which have been voiced throughout the philosophical community for some time—made it easier to understand why a modificationist might look for non-relational alternatives to the relation view of intentionality.

Next I examined three views within the phenomenal intentionalist camp—focusing in on two similar views, and distinguishing them as modificationist. I explained the elements of modificationism, and went on to address some concerns that arise for phenomenal intentionalism in general and modificationism in particular. I noted several different attempts to address these concerns in the phenomenal intentionality literature, and suggested that some promising progress has been made.

I then moved on to what I take to be a rather large concern for modificationism—one that I have heard described as the elephant in the phenomenal intentionalist room: the problem of cognitive contact. After explaining what the problem is, and why it emerges as particularly vexing given the theoretical constraints of modificationism, I went on to examine two theories that served as exemplars for two different approaches to cognitive contact. The theories were sense data theory and disjunctivism, and their respective approaches were what I called indirectivism and directivism, respectively. From my examination of directivism and indirectivism, I concluded that a perfect theory of cognitive contact is hard to come by, and therefore that we might have to settle for a less than perfect account of how we manage to get outside our heads. Keeping these more modest expectations in mind, I went on to examine the options a modificationist might have to address the problem of cognitive contact. Though in need of a great deal more theorizing, I suggested two avenues down which a modificationist might pursue a solution to the problem of cognitive contact. My conclusion was that modificationism is not necessarily doomed with respect to providing a solution to the problem of cognitive contact.

In summary, my aspirations were modest: I wanted to investigate what I took to be an interesting theory, highlight some of its central motivations, and examine its strengths and weaknesses—the most problematic of which is modificationism's seeming inability to account for our cognitive contact with the world outside our skins. Whether or not phenomenal intentionality theory, and modificationism, can deliver a viable account of cognitive contact remains to be seen. All the same, I hope this project serves, at the very least to draw attention to the problem; and, at best, gives hope that a solution to the problem of cognitive contact, even an uncompromising one, might be possible.

Bibliography

Adams, F., & Aizawa, K. (1992). 'X' means X: Semantics Fodor-Style. *Minds and Machines , 2*, 175-183.

Anscombe, G. (2002). The Intentionality of Sensation: A Grammatical Feature. In A. Noe, & E. (. Thompson (Eds.), *Vision and Mind* (pp. 155-175). Cambridge, Massachusets, USA: The MIT Press.

Armstrong, D. (1997). Against 'Ostrich' Nominalism: A Reply to Michael Devitt. In D. Mellor, & A. Oliver (Eds.), *Properties* (pp. 101-111). New York, New York, USA: OUP.

Armstrong, D. (1997b). Properties. In D. Mellor, & A. Oliver (Eds.), *Properties* (pp. 160-172). New York, New York, USA: OUP.

Armstrong, D. (1999). Universals as Attributes. In J. Kim, & E. Sosa (Eds.), *Metaphysics: An Anthology* (pp. 198-208). Malden, Massachusetts, USA: Blackwell Publishing Ltd.

Bayne, T., & Montague, M. (Eds.). (2011). *Cognitive Phenomenology.* New York, New York, USA: OUP.

Bickhard, Mark H. (2010) Interactive Knowing: The Metaphysics of Intentionality. In Poli, R. & Sebibt, J (Eds.) *Theory and Applications of Ontology: Philosophical Perspectives.* New York, New York, USA: Springer

Bilgrami, A. (1996). Can Externalism Be Reconciled With Self-Knowledge? In A. Pessin, & S. Goldberg (Eds.), *The Twin Earth Chronicles: Twenty Years of Reflection on Hilary Putnam's "The Meaning of 'Meaing'"* (pp. 362-394). Armonk, New York, USA: M.E. Sharpe Inc.

Block, N. (1995). On a Confusion about the Function of Consciousness. *Behavioral and Brain Sciences* (18), 227-247.

Block, N. (Ed.). (1981). *Readings in the Philosophy of Psychology* (Vol. 2). Cambridge, Massachusetts, USA: Harvard University Press.

Boghossian, P. A. (1998). What the Externalist Can Know "A Priori". *Philosophical Issues , 9*, 197-211.

Boghossian, P. A., & Velleman, D. J. (1997). Colour as a Secondary Quality. In *Readings on Color: The Philosophy of Color* (Vol. 1, pp. 81-104). Cambridge, Massachusetts, USA: The MIT Press.

Bourget, D. (2010). Consciousness is Underived Intentionality. *Nous , 44* (1), 32-58.

Brentano, F. (1887/1995). *Descriptive Psychology.* (B. Muller, Trans.) New York, New York, USA: Routledge.

Burge, T. (1996). Individualism and Self-Knowledge. In A. Pessin, & S. Goldberg (Eds.), *The Twin Earth Chronicles: Twenty Years of Reflection on Hilary Putnam's "The Meaning of 'Meaing'"* (pp. 342-354). Armonk, New York, USA: M.E. Sharpe.

Burge, T. (1979). Individualism and the Mental. *Midwest Studies in Philosophy , 4* (1), 73-122.

Burge, T. (1991). Vision and Intentional Content. In E. Lepore, & R. Van Guilick (Eds.), *John Searle and His Critics* (pp. 195 - 214). Cambridge, Massachusetts, USA: Blackwell Publishers.

Byrne, A., & Hilbert, D. (2003). Color Realism and Color Science. *Behavioral and Brain Sciences* (37).

Byrne, A., & Logue, H. (2009). Introduction. In A. Byrne, & H. Logue (Eds.), *Disjunctivism* (pp. vii-xxix). Cambridge , Massachusetts, USA: The MIT Press.

Campbell, K. (1997). The Metaphysic of Abstract Particulars. In D. Mellor, & A.

Oliver (Eds.), *Properties* (pp. 125-139). New York, New York, USA: OUP.

Chalmers, D. J. (Ed.). (2002). *Philosophy of Mind: Classical and Contemporary Readings.* New York, New York, USA: OUP.

Chisholm, R. (1957). *Perceiving: A Philosophical Study.* Ithica, New York, USA: Cornell University Press.

Churchland, P. (2005). Chimerical colors: some phenomenological predictions from cognitive neuroscience. *Philosophical Psychology , 18* (5), 527-560.

Coval, S., & Campbell, P. (2010). *Critique of the Liberal Idea of a Person.* Lewiston, New York, USA: The Edwin Mellen Press.

Crane, T. (1996). All the Difference in the World. In *The Twin Earth Chronicles* (pp. 284-304). New York, New York, USA: M.E. Sharpe.

Crane, T. (2001). Intentional Objects. *Ratio , XIV* (4), 336-349.

Daly, C. (1997). Tropes. In D. Mellor, & A. Oliver (Eds.), *Properties* (pp. 140-159). New York, New York, USA: OUP.

Dancy, J. (2009). Arguments from Illusion. In A. Byrne, & H. Logue (Eds.), *Disjunctivism* (pp. 117-136). Cambridge , Massachusetts, USA: The MIT Press.

Davidson, D. (1967/2001). *Inquiries into Truth and Interpretation.* New York, New York, USA: OUP.

Davidson, D. (1996). Knowing One's Own Mind. In A. Pessin, & S. Goldberg (Eds.), *The Twin Earth Chronicles: Twenty Years of Reflection on Hilary Putnam's "The Meaning of 'Meaing'"* (pp. 323-341). Armonk, New York, USA: M.E. Sharpe.

Dennett, D. C. (1987). *The Intentional Stance.* Cambridge, Massachusets, USA: The MIT Press.

Dennett, D. C. (1969/1986). *Content & Consciousness.* New York, New York, USA: Routledge & Kegan Paul plc.

Devitt, M. (1997). 'Ostrich Nominalism' or 'Mirage Realism'. In D. Mellor, & A. Oliver (Eds.), *Properties* (pp. 93-100). New York, New York, USA: OUP.

Donnellan, K. (2008). Reference and Definite Descriptions. In *The Philosophy of Language* (pp. 265-277). New York/Oxford, New York, USA/UK: OUP.

Dretske, F. (1981). *Knowledge and the Flow of Information.* Cambridge, Massachusetts, USA: The MIT Press.

Dretske, F. (1997). *Naturalizing the Mind.* Cambridge, Massachusetts, USA: The MIT Press.

Dreyfus, H. L., & Hall, H. (Eds.). (1984). *Husserl, Intentionality and Cognitive Science.* Cambridge , Massachusets, USA: The MIT Press.

Ducasse, C. (1942). Moore's Refutation of Idealism. In P. Schlipp (Ed.), *The Philosophy of G.E. Moore* (pp. 223-252). Chicago, Illinois, USA: Northwestern University Press.

Dummett, M. (1991). *Frege and Other Philosophers.* New York, New York, USA: OUP.

Dummett, M. (2006). *Thought and Reality.* New York, New York, USA: OUP.
Farkas, K. (2008). Phenomenal Intentionality without Compromise. *Monist ,*
91 (2), 273-293.

Fodor, J. (1990). *A Theory of Content.* Cambridge , Massachusetts, USA: The
MIT Press.

Fodor, J. (1998). *Concepts.* New York, New York, USA: OUP.

Fodor, J. (2008). *LOT 2 The Language of Thought Revisited.* New York, New
York, USA: OUP.

Fodor, J. (1987). *Psychosemantics.* Cambridge , Massachusetts, USA: The MIT
Press.

Fodor, J. (1975). *The Language of Thought.* Cambridge, Massachusetts, USA:
Harvard University Press.

Fodor, J., & Lepore, E. (1992). *Holism: A Shopper's Guide.* Cambridge,
Massachusets, USA: Blackwell Publishers.

Follesdal, D. (1984). Husserl's Notion of Noema. In H. L. Dreyfus, & H. Hall
(Eds.), *Husserl, Intentionality and Cognitive Science* (pp. 73 - 80). Cambridge,
Massachusetts, USA: The MIT Press.

Frege, G. (1892/1997). Comments on Sinn and Bedeutung. In *The Frege*
Reader (pp. 172-180). Malden, Massachusetts, USA: Blackwell Publishing Ltd.

Frege, G. (1891/1997). Function and Concept. In D. Mellor, & A. Oliver (Eds.),
Properties (pp. 34-44). New York, New York, USA: OUP.

Frege, G. (1892/1997). On Concept and Object. In M. Beaney (Ed.), *The Frege*
Reader (pp. 181-193). Malden, Massachusetts, USA: Blackwell Publishing Ltd.

Frege, G. (1892/1997). On Sinn and Bedeutung. In M. Beaney (Ed.), *The Frege*
Reader (pp. 151-171). Malden, Massachussets, USA: Blackwell Publishing Ltd.

Frege, G. (1918/1997). Thought. In M. Beaney (Ed.), *The Frege Reader* (pp.
325-345). Malden, Massachusetts, USA: Blackwell Publishing LTD.

Gendler, T. S., & Hawthorne, J. (Eds.). (2006). *Perceptual Experience.* New
York, New York, USA: OUP.

Georgalis, N. (2006). *The Primacy of the Subjective.* Cambridge, Massachusetts, USA: The MIT Press.

Gurwitsch, A. (1984). Husserl's Theory of Intentionality and Conscious. In H. L. Dreyfus, & H. Hall (Eds.), *Husserl, Intentionality and Cognitive Science* (pp. 59-72). Cambridge , Massachusetts, USA: The MIT Press.

Hardin, C. (1988). *Color for Philosophers.* Indianapolis/Cambridge, Indiana/Massachusetts, USA: Hacket Publishing Company.

Hawley, K., & Macpherson, F. (Eds.). (2011). *The Admissible Contents of Experience.* Malden, Massachusetts, USA: Balckwell.

Horgan, T., & Graham, G. (2012). Phenomenal Intentionality and Content Determinacy. In R. Shantz (Ed.), *Prospects for Meaning* (pp. 321-344). Berlin/Boston, Massachusetts, Germany/USA: De Gruyter.

Horgan, T., & Tienson, J. (2002). The Intentionality of Phenomenology and the Phenomenology of Intentionality. In D. J. Chalmers (Ed.), *Philosophy of Mind: Classical and Contemporary Readings.* New York, New York, USA: OUP.

Horgan, T., Tienson, J., & Graham, G. (2004). Phenomenal Intentionality and the Brain in a Vat. In R. Shantz (Ed.), *The Externalist Challenge* (pp. 297 - 317). New York, New York, USA: Walter de Gruyter GmbH & Co.

Husserl, E. (1913/1998). *Ideas Pertaining to a Pure Phenomenology and to a Phenomenological Philosophy* (Vol. I). (F. Kersten, Trans.) Boston, Massachusetts, USA: Kluwer Academic Publishers .

Husserl, E. (1900/2001). *Logical Investigations* (Vol. 1). New York, New York, USA: Routledge.

Jackson, F. (1977). *Perception: A Representative Theory.* New York, New York, USA: CUP.

Johnson, M. (2009). The Obscure Object of Hallucination. In A. Byrne, & H. Logue (Eds.), *Disjunctivism* (pp. 216-269). Cambridge, Massachusetts , USA: The MIT Press.

King, J. C. (2007). *The Nature and Structure of Content.* New York, New York, USA: OUP.

Klein, C. (2007). An Imperative Theory of Pain. *Journal of Philosophy , 104* (10), 517-532.

Kriegel, U. (2011). Cognitive Phenomenology as the Basis of Unconscious Content. In T. Bayne, & M. Montague (Eds.), *Cognitive Phenomenology* (pp. 79-102). New York, New York, USA: OUP.

Kriegel, U. (2007). Intentional Inexistence and Phenomenal Intentionality. *Philosophical Perspectives , 21* (1), 307-340.

Kriegel, U. (2013). *Phenomenal Intentionality.* (U. Kriegel, Ed.) New York, New York, USA: OUP.

Kriegel, U. (2009). *Subjective Consciousness.* New York, New York, USA: OUP.

Kriegel (Ed.), *Phenomenal Intentionality.* New York, New York, USA: OUP.

Kriegel, U. (2011). *The Sources of Intentionality.* New York, New York, USA: Oxford University Press.

Kriegel, U. (2011b). The Veil of Abstracta. *Philosophical Issues* (21).

Kripke, S. (1972). *Naming and Necessity.* Cambridge, Massachusetts, USA: Harvard University Press.

Kripke, S. (1982). *Wittgenstein On Rules and Private Language.* Cambridge, Massachusetts, USA: Harvard University Press.

Levine, J. (2011). On the Phenomenology of Thought. In T. Bayne, & M.

Montague (Eds.), *Cognitive Phenomenology* (pp. 103-120). New York, New York, USA: OUP.

Lewis, D. (1997). Modal Realism at Work: Properties. In D. Mellor, & A. Oliver (Eds.), *Properties* (pp. 160-172). New York, New York, USA: OUP.

Lewis, D. (1997). New Work for a Theory of Universals. In D. Mellor, & A. Oliver (Eds.), *Properties* (pp. 188-227). New York, New York, USA: OUP.

Lewis, D. (1986). *On the Plurality of Worlds.* Malden , Massachusetts, USA: Blackwell Publishing.

Loar, B. (1981). *Mind and Meaning.* New York, New York, USA : Cambridge University Press.

Loar, B. (2003). Phenomenal Intentionality as the Basis of Mental Content. In M. Hahn, & B. Ramberg (Eds.), *Reflections and Replies: Essays on the Philosophy of Tyler Burge* (pp. 229-257). Cambridge, Massachusetts, USA: The MIT Press.

Loar, B. (1987). Subjective Intentionality. *Philosophical Topics , 15* (1), 89-124.

Lycan, W. (1995). *Consciousness.* Cambridge , Massachusetts, USA: The MIT Press.

Lycan, W. G. (1996). *Consciousness and Experience.* Cambridge, Massachusetts, USA: The MIT Press.

Makin, G. (2000). *The Metaphysicians of Meaning: Russell and Frege on sense and denotation.* New York, New York, USA: Routledge.

Martin, M. (2009). The Reality of Appearances. In A. Byrne, & H. Logue (Eds.), *Disjunctivism* (pp. 91-115). Cambridge, Massachusetts, USA: The MIT Press.

McKinsey, M. (1996). Anti-Individualism and Priviledged Access. In A. Pessin, & S. Goldberg (Eds.), *The Twin Earth Chronicles: Twenty Years of Reflection on Hilary Putnam's "The Meaning of 'Meaing'"* (pp. 355-361). Armonk, New York, USA: M.E. Sharpe Inc.

Meinong, A. (1960). On the Theory of Objects. In R. Chisholm (Ed.), *Realism and the Background of Phenomenology* (pp. 76-117). Glencoe, Illinois, USA: Free Press.

Mendelovici, A. (2010, September). Mental Representation and Closely Conflated Topics. *PhD Thesis* . Princeton, New Jersey, USA: Princeton University.

Mendelovici, A. (2013, March). Mental Representation is not a Relation.

Mendelovici, A. (MS). Phenomenal Intentionality: How to Get Intentionality from Phenomenal Consciousness.

Millikan, R. G. (2002). Biosemantics. In *Philosophy of Mind: Classic and Contemporary Readings* (pp. 500-508). New York, New York, USA: OUP.

Millikan, R. G. (2000). *On Clear and Confused Ideas.* New York, New York, USA: Cambridge University Press.

Moltmann, F. (2014). Propositions, attitudinal objects and the distinction between actions and products. *Canadian Journal of Philosophy* .

Montague, M. (2007). Against Propositionalism. *Nous , 41* (3), 503-518.

Montague, M. (2013). The Access Problem. In U. Kriegel (Ed.), *Phenomenal Intentionality* (pp. 27-48). New York, New York, USA: OUP.

Nagel, T. (1980). What Is It Like to Be a Bat? In N. Block (Ed.), *Readings in the Philosophy of Psychology* (Vol. 1, pp. 159-168). Cambridge, Massachusetts, USA: Harvard University Press.

Noe, A. (2004). *Action in Perception.* Cambridge, Massachusetts, USA: The MIT Press.

Parsons, T. (1975). A Meinongian Analysis of Fictional Objects. *Grazer Philosophische Studien , 1*, 73-86.

Pessin, A., & Goldberg, S. (Eds.). (1996). *The Twin Earth Chronicles: Twenty Years of Reflection on Hilary Putnam's "The Meaning of 'Meaning' "".* Armonk, New York, USA: M.E. Sharpe Inc.

Pietroski, P. (1992). Intentional and Teleological Error. *Pacific Philosophica Quarterly , 73* (3), 267 - 282.

Pitt, D. (2009). Intentional Psychologism. *Philosophical Studies , 146* (1), 117-138.

Pitt, D. (2011). Introspection, Phenomenality, and the Availability of Intentional Content. In T. Bayne, & M. Montague (Eds.), *Cognitive Phenomenology* (pp. 141-173). New York, New York, USA: OUP.

Pitt, D. (2013). *Mental Representation.* From The Stanford Encyclopedia of Philosophy: <http://plato.stanford.edu/archives/fall2013/entries/mental-representation/>

Pitt, D. (2004). The Phenomenology of Cognition OR What It Is Like To Think That P. *Philosophy and Phenomenological Research , LXIX* (1).

Prinz, J. (2007). All Consciousness is Perceptual. In B. McLaughlin, & J. Cohen (Eds.), *Contemporary Debates in Philosophy of Mind* (pp. 335-357). Malden, Massachusetts, USA: Blackwell Publishing Ltd.

Prinz, J. (2011). The Sensory Basis of Cognitive Phenomenology. In T. Bayne, & M. Montague (Eds.), *Cognitive Phenomenology* (pp. 174-196). New York, New York, USA: OUP.

Putnam, H. (1975). The Meaning of "Meaning". In K. Gunderson (Ed.), *Language, Mind and Knowledge* (Vol. VII, pp. 131-193). Minneapolis, Minnesota, USA: University of Minnesota Press.

Putnam, H. (1996). The Meaning of 'Meaning'. In A. Pessin, & S. Goldberg (Eds.), *The Twin Earth Chronicles: Twenty Years of Reflection on Hilary Putnam's "The Meaning of 'Meaing'"* (pp. 3-52). Armonk, New York, USA: M.E. Sharpe.

Putnam, H. (1975). The Meaning of 'Meaning". *Minesota Studies in the Philosophy of Science* (7), 131-193.

Quine, W. (1953). *From A Logical Point of View.* Cambridge, Massachusets, USA: The Harvard University Press.

Quine, W. (1966). Russell's Ontological Development. *The Journal of Philosophy , 63* (21), 657-667.

Quine, W. (1966). *The Ways of Paradox and Other Essays.* Cambridge, Massachusetts, USA: Harvard University Press.

Quine, W. (1960). *Word and Object.* Cambridge, Massachusetts, USA: The MIT Press.

Ramsey, F. (1997). Universals. In D. Mellor, & A. Oliver (Eds.), *Properties* (pp. 57-73). New York, New York, USA: OUP.

Reid, T. (1983). *Inquiry and Essays.* (R. Beanblossom, & K. Lehrer, Eds.) Indianapolis, Indiana, USA: Hacket Publishing Company, Inc.

Russell, B. (1905/2008). On Denoting. In A. Martinich (Ed.), *The Philosophy of Language* (pp. 230-238). New York, New York, USA: OUP.

Russell, B. (1904/1988). Selection From the Frege-Russell Correspondence. In N. Salmon, & S. Soames, *Propositions and Attitudes* (p. 57). New York, New York, USA: OUP.

Russell, B. (1927). *The Analysis of Matter.* New York, New York, USA: Harcourt, Brace.

Russell, B. (1903/1996). *The Principles of Mathematics.* New York, New York, USA: W.W. Norton & Company.

Russell, B. (1912/1997). *The Problems of Philosophy .* New York, New York, USA: OUP .

Russell, B. (1967/1997). The World of Universals. In D. Mellor, & A. Oliver (Eds.), *Properties* (pp. 45-50). New York, New York, USA: OUP.
Russell, B., & MacColl, H. (1905). The Existential Import of Propositions. *Mind , 14* (55), 398-402.

Schitzgebel, E. (2012). Introspection, What? In D. Smithies, & D. (. Stoljar (Eds.), *Introspection And Consciousness* (pp. 29-47). New York, New York, USA: OUP.

Searle, J. (1990). Consciousness, Explanatory Inversion and Cognitive Science. *Behavioral and Brain Sciences , 13*, 585-642.

Searle, J. (1987). Indeterminacy, Empiricism, and the First Person. *Journal of Philosophy , 84* (3), 123-146.

Searle, J. (1983). *Intentionality: An Essay in the Philosophy of Mind.* Cambridge, UK: CUP.

Searle, J. (1958). Proper Names. *Mind , 67* (266), 166-173.

Searle, J. (1992). *The Rediscovery of the Mind.* Cambridge, Massachusetts, USA: The MIT Press.

Segal, G. (2000). *A Slim Book About Narrow Content.* Cambridge, Massachusetts, USA: The MIT Press.

Segal, G. (2004). Reference, Causal Powers, Externalist Intuitions and Unicorns. In R. Shantz (Ed.), *The Externalist Challenge* (pp. 329-344). New York, New York, USA: Walter de Gruyter GmbH & Co.

Shantz, R. (Ed.). (2004). *The Externalist Challenge* (Vol. 2). New York, New York, USA: Walter de Gruyter GmbH & Co.

Shields, C. (2011). On Behalf of Cognitive Qualia. In T. Bayne, & M. Montague (Eds.), *Cognitive Phenomenology* (pp. 215-235). New York, New York, USA: OUP.

Shoemaker, S. (1997). Causality and Properties. In *Properties* (pp. 228-254). New York, New York, USA: OUP.

Siegel, S. (2011). *The Contents of Visual Experience.* New York, New York, USA: OUP.

Siegel, S. (2011). The Visual Experience of Causation. In K. Hawley, & F. Macpherson (Eds.), *The Admissible Contents of Experience.* Malden, MA, USA: Blackwell.

Siewert, C. P. (1998). *The Significance of Consciousness.* Princeton, New Jersey, USA: Princeton University Press.

Siewert, C. (2011). Phenomenal Thought. In T. Bayne, & M. Montague (Eds.), *Cognitive Phenomenology* (pp. 236-267). New York, New York, USA: UOP.

Smithies, D., & Stoljar, D. (Eds.). (2012). *Introspection and Consciousness.* New York, New York, USA: OUP.

Soames, S. (2014). The Inadequacy of the Traditional Conceptions. In J. C. King, S. Soames, & J. Speaks, *New Thinking About Propositions.* New York, New York, USA: OUP.

Soames, S. (2010). *What is Meaning?* Princeton, New Jersey, USA: Princeton University Press.

Stalnaker, R. (1996). On What's in the Head. In *The Twin Earth Chronicles* (pp. 238-263). New York, New York, USA: M.E. Sharpe.

Stampe, D. (1977). Toward a Causal Theory of Linguistic Representation. *Midwest Studies in Philosophy , 2* (1), 42-63.

Strawson, G. (2011). Cognitive Phenomenology: Real Life. In T. Bayne, & M. Montague (Eds.), *Cognitive Phenomenology* (pp. 285-325). New York, New York, USA: OUP.

Strawson, G. (2010). *Mental Reality* (2nd ed.). Cambridge, Massachusetts, USA: The MIT Press.

Strawson, G. (2004). Real Intentionality. *Phenomenology and the Cognitive Science , 3* (3), 287-313.

Strawson, G. (2008). *Real Materialism.* New York, New York, USA: OUP.

Tye, M. (1995). A Representational Theory of Pains and their Phenomenal Character. *Philosophical Perspectives , 9*, 223-239.

Tye, M. (2009). *Consciousness Revisited.* Cambridge, Massachusets, USA: The MIT Press.

Tye, M. (2002). *Consciousness, Color, and Content.* Cambridge, Massachusetts, USA: The MIT Press.

Tye, M. (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind.* Cambridge, Massachusetts, USA: The MIT Press.

Tye, M., & Wright, B. (2011). Is There a Phenomenology of Thought? In T. Bayne, & M. Montague (Eds.), *Cognitive Phenomenology* (pp. 326-344). New York, New York, USA: OUP.

Van Cleve, J. (1985). Three Versions of the Bundle Theory. *Philosophical Sutdies , 47* (1), 95-107.

Williams, D. C. (1997). On the Elements of Being: I. In D. Mellor, & A. Oliver (Eds.), *Properties* (pp. 112-124). New York, New York, USA: OUP.

Wittgenstein, L. (1922/1961). *Tractatus Logico-Philosophicus.* (D. Pears, & B. McGuiness, Trans.) New York, New York, USA: Routledge.

Yablo, S. (2010). *Things: Philosophical Papers Volume II.* New York, New York, USA: OUP.

Yablo, S. (2008). *Thoughts: Philosophical Papers Volume I.* New York, New York, USA: OUP.

CV

C.A. Young

Post-secondary Education and Degrees:
University of Regina
Regina, Saskatchewan, Canada
2006-2009

Honours and Awards:

Ontario Graduate Scholarship
2011-2012

Related Work Experience:

Teaching Assistant
The University of Western Ontario
2009-2013