RICE UNIVERSITY

Perceptual Links:
Attention, Experience, and Demonstrative Thought

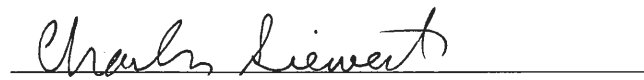by

Michael Barkasi

A THESIS SUBMITTED
IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE

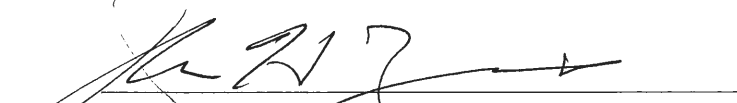## Doctor of Philosophy

APPROVED, THESIS COMMITTEE:

Casey O'Callaghan, Associate Prof. of Philosophy,
Washington University in St. Louis, Thesis Advisor

Charles Siewert, Robert Alan and Kathryn Dunlevie
Hayes Prof. of Humanities and Prof. of Philosophy

Richard Grandy, McManis Prof. of Philosophy,
Committee Chair

John H. Zammito, John Antony Weir Prof.
of History

ABSTRACT

Perceptual Links:
Attention, Experience, and Demonstrative Thought

by

Michael Barkasi

Perception is conscious: perceiving involves a first-person experience of what's perceived. It's widely held that these perceptual experiences are independent of what's perceived. Viewing two visually indiscriminable #2 pencils would involve the same experience, despite viewing different objects. It's also widely held that conscious perception enables thinking about what's perceiving. When you see one of those pencils you can think, *that* is a pencil. Some philosophers, including John McDowell and John Campbell, have suggested that these two features engender a puzzle: how can perceptual experiences make perceived objects available for thought when they're independent of those objects? This dissertation is a collection of four papers which address this question.

The first (chapter 2) argues that, under two minimal assumptions, conscious perception makes objects available for thought only if experience is not object independent. The second (chapter 3) argues that conscious perception makes objects available for thought by enabling voluntarily attention to them. The third (chapter 4) integrates empirical work on multiple-object tracking and philosophical work on attention to argue that conscious perception isn't mediated by the construction of representations within the visual system. The fourth (chapter 5) uses philosophical methods and neurophysiology to give an account of the role of experience itself in how perception makes objects available for thought. A concluding chapter combines and extends results from the previous chapters to give a naïve realist (vs representationalist) account of perceptual experience.

The questions about perceptual experience addressed in this dissertation (object dependent or independent? naïve realist or representationalist?) are fundamental to our understanding of experience. Not only do they get at its basic nature, but their answers constrain how we might give scientifically respectable, or naturalistic, explanations of experience as well as how we might explain perceptual hallucinations and illusions.

Acknowledgments

This project started when Casey O'Callaghan suggested I turn a paper I had written for him into my dissertation. This was during the spring of 2012, while I was doing an independent reading course with him as part of my preparation for qualifying exams. John Campbell and Bill Brewer had each given controversial arguments for naïve realism, arguing that it was required to explain perception's role in demonstrative thought. My paper suggested that within these more controversial arguments was a more straightforward one for the object dependence of experience. That paper lead to another one, also drawing on earlier work from John McDowell, which I gave the following May (2013) in Riga at the 9th International Symposium of Cognition, Logic and Communication. By that point I had been struggling for a year to fill serious gaps in the argument and to go meaningfully beyond Campbell, McDowell, and Brewer. Later that same summer, while back in Riga as a student at ISSCSS (the International Summer School in Cognitive Sciences and Semantics), a number of points finally became clear. There I put together a presentation in which I gave a relatively gap-free argument for object dependence from demonstrative thought. Unlike my previous papers the presentation wasn't just an explication of Campbell, McDowell, and Brewer's work, but developed my own original argument. That presentation lead to a paper, which eventually became chapter 2. That paper, written in fall 2013, proved to be a fecund starting point. Over the next academic year (2013–14) and summer (2014) I quickly wrote three more papers (chapters 3–5) and the material for the conclusion (chapter 6).

I have to thank Casey for encouraging me to pursue this project. On my own I wouldn't have seen its potential. More importantly, I have to thank him for his willingness to take on the project with me as my advisor. This, of course, involved reading many drafts and spending a lot of time discussing my ideas. Over the last few years Casey also helped me develop a sense for philosophical writing and served as a model for doing philosophy well. Whatever success I achieve here is due mostly to his help and guidance over that time. Going back even further,

Casey helped me stay on track in my early days at Rice and always was a source of sound guidance.

A number of other people provided valuable input. Richard Grandy and Indrek Reiland deserve special mention. Richard has probably discussed and provided input on nearly every major draft of chapters 2–4, usually at a moment's notice. Indrek was with me in Riga for ISSCSS when I made my breakthrough on the argument for object dependence. He provided a week's worth of discussion of the issues which helped me get clear on a number of points. He also later provided extensive written comments on later drafts of chapter 2 and early drafts of chapters 3 and 4. I've also had several productive conversations with Charles Siewert about the material from these three chapters. James Genone and Jeff Speaks provided helpful comments on early drafts of chapter 2. Brian McLaughlin provided helpful comments on chapter 2 as my commentator at the 2014 APA Pacific. Dan Burnston, Alex Morgan, and Roma Hernández provided helpful thoughts on the material from chapters 3–5 in a series of conversations over the 2014–2015 academic year. The audience at the 2013 Riga Symposium gave helpful comments on the intermediate paper that lead to chapter 2. Audiences at the 2013 ISSCSS and 2014 APA Pacific gave helpful comments on the material from chapter 2. At the Pacific Imogen Dickie gave me a few suggestions which partly prompted a complete rewrite of the chapter. I received helpful comments on both chapter 2 and chapter 3 during work-in-progress talks at Rice (April 17 and October 25, 2013). I also received several sets of helpful anonymous reviews on drafts of chapters 2, 3, and 4 from journal rejections.

In addition to helpful input, I need to thank those who helped in practical ways. All my committee members need to be thanked: Casey, Richard, Charles, and John Zammito. Those who attended my dissertation defense on April 17th were also crucial. After all, I could not have *passed* my defense without a *defense*. All in attendance asked great questions which will help in any future work derived from the dissertation.

Of course, successfully writing a dissertation and making it through a doctoral program requires layers of personal and professional support. The Rice Philosophy Department has provided me not only with the usual financial support but also with more-than-generous support for attending international summer schools and conferences. During my time in the Rice department Richard Grandy has especially gone above in providing mentorship, advice, and support. I've benefited deeply from many dear friends in Houston, especially Jake Mills, Chris Biehle, Maria Picone, Lorena Maili, Sean De Vega, Joe Adams, and Roma

Hernández. Finally, I would not have made it through the last few years of graduate school were it not for the constant and unwavering support of my wife, Coretta Gerould.

As the final drafts of all the chapters were being prepared concurrently with the introduction and the introduction is a survey of work in the rest of the dissertation, some of the material in the introduction is an adaptation of material in later chapters. In a few cases bits of papers are borrowed word-for-word in the introduction. Specifically, §1.2.1 and §1.2.2 of the introduction are heavy extended excerpts from §2.2. §1.4.1 contains extended excerpts from §2.3 and §2.4. The second and third paragraphs of §1.3.2 are a slightly modified excerpt from §5.2. The second paragraph of §1.2.4 is a modified except from §4.3.

Contents

List of Figures

List of Tables

Chapter 1:
Introduction

1.1 Overview

Conscious perception introduces what's perceived for thought. Specifically, it allows you to identify or select it for thought. For example, look around and focus on a nearby pencil. You might think, *that* is a #2 pencil. But a moment's reflection raises a question (McDowell 1986, 165; Martin 2002a, 198). If the pencil was instantaneously swapped for another that's visually indiscriminable, your experience would remain the same. Similarly, there might be no pencil there at all, but you hallucinate in a way that your experience is the same as it would be were the pencil actually there.

1. How could conscious perception of the pencil allow you to select it for thought when, as these two examples suggest, your experience is independent of that particular pencil?

The question suggests two others:

2. How does conscious perception allow you to select things for thought?
3. What role, if any, is played by experience?

My dissertation is a collection of four independent papers (chapters 2–5) and a final chapter (chapter 6) which address these questions.

The first paper (chapter 2) addresses question 1. It argues that if

(a) conscious perception allows you to select what's perceived by enabling demonstrative selection of it, and
(b) experience plays some role,

then contrary to what the above two examples suggest, experience is not object independent. This is interesting because (a) and (b) are fairly noncontroversial responses to questions 2 and 3, respectively. So, the first paper shows that the

1

worry implicit in question 1 is well-founded: on a fairly noncontroversial picture of how conscious perception introduces what's perceived for thought, experience is object dependent. In addition to being unintuitive, object-dependence entails that two widely held views are false (see §1.4.1). The first goes back to Descartes and is that features of experience are fully accessible through introspection. The second is that experiences are fully determined by what goes on within the biological systems of the perceiver ("in the head").

The next three papers (chapters 3–5) address questions 2 and 3. A plausible view is that in demonstrative selection you select a target of thought by exploiting a relation between you and it. So if conscious perception enables demonstrative selection, what specifically is the exploited relation? What is, as Imogen Dickie calls it (2010, 213), the *perceptual link*? Voluntary focal attention to consciously perceived things relates you to them. Chapter 3 argues that it's this relation that's exploited in perception-based thought. Chapter 4 argues that the relation of voluntary focal attention—as well as conscious perceiving itself—isn't mediated by representations in the visual system, a result crucial for chapter 6. Chapter 5 moves on to question 3 and argues that experience's role in perception-based thought is to provide perceptual information for use in the voluntary control of focal attention.

Chapters 3–5 add up to an account of the central aspects of perception-based thought: how perception allows for it and what role experience plays. Having an account is important because perception-based thought plays a fundamental role in our cognitive lives (e.g., Campbell 1997, 55; Pylyshyn 2007, 19; Levine 2010, 169). Although I don't discuss them in the dissertation, it's worth mentioning three specific roles perception-based thought is often taken to play. First, it's a way to have nondescriptive, relational thought (Campbell 2002; Raftopoulos and Müller 2006, 253,261–2; Bach 2010, 55; Jeshion 2010c, 134; Recanati 2012, 12,37; Dickie forthcoming). Second, it's a way of selecting targets for thought without having concepts of them (Campbell 2002; Raftopoulos and Müller 2006, 252–3; Pylyshyn 2007, 56). Finally, it's needed for acquiring concepts through experience (Putnam 1975; Raftopoulos and Müller 2006, 261; cf. Levine 2010, 193). Learning the concept RED, for example, requires seeing instances of red and thinking, *that* is red.

Chapter 6 uses the result from chapter 4 on representations in the visual system to extend the argument from chapter 2 into an argument for naïve realism. Naïve realism says that experiences are relational. It's main alternative—and also the dominant view—is representationalism, the view that experiences are repre-

sentations. Although naïve realism entails that experiences are object dependent, there are versions of representationalism on which experiences are also object dependent. The work of chapter 6 is to rule these out. Whether experience is relational or representational is a fundamental question in philosophy of perception (Crane 2006). First, the debate concerns the nature of perceptual experience—is it relational, or representational?—and perceptual experience is a central topic of the field. Second, if experience is relational then a common strategy for naturalizing experiences is ruled out. This strategy (e.g., Tye 1995; Dretske 2003; see Bourget and Mendelovici 2014 for overview) is to first show that experiences are representations, then give a naturalized account of representations such as Dretske, Millikan, and Fodor's information-based accounts (Dretske 1981, 1988, 1995; Millikan 1984, 1989, 2004, 2009; Fodor 1987, 1990). Third, whether experience is relational or representational constrains explanations of illusion and hallucination (and questions about illusion and hallucination are also central to the field).

The puzzle with which I started (question 1) is one version of a deep problem that's historically faced empiricism. Empiricism says the experiential aspects of conscious perception introduce some things for thought (e.g., Hume 1993/1777, §2). But experience seems to have features which make this impossible. The version I give, which takes object independence as the problematic putative feature, comes from John McDowell (1986, 165; see also Pettit and McDowell 1986, 13–15). Like me, McDowell argues from it that experience is not object independent. More recently John Campbell (2002, 125, 2010, 201), Bill Brewer (2006, 179), and Mark Johnston (2006, 264) have used this version to argue for naïve realism. Campbell and Brewer trace the puzzle back to Berkeley, who focuses on a different putative feature of experience (Brewer 2011, 18–21,36–41; Roessler 2011, 2). Locke believed that experiences present ideas caused by things in the world, not the things themselves. Berkeley saw that this was in tension with his view that the experiential aspects of conscious perception introduce some things for thought (see also Russell 1956, 147). Of course, Berkeley responded by identifying the world with the ideas, not by rejecting the claim that experiences present ideas.

My dissertation goes beyond the work from McDowell, Campbell, Brewer, and Johnston in several ways. They largely fail to develop the intuitive force of question 1 into an explicit argument for the object dependence of experience (Burge 2005, 60,62; Pautz 2010, 283,286). I develop an argument (chapter 2) using (a) and (b) as premises. This argument has two innovative features. First,

the argument leverages the relational character of demonstrative thought in an insightful way. The following outline of the argument brings out this leveraging. If (a) conscious perception enables demonstrative selection, then it provides an exploitable relation for demonstrative selection; for (b) experience to play a role, it must be involved in that relation, making it object dependent. The second innovative feature is that the argument fills out this outline by giving two conditions jointly sufficient for object dependence, then showing that one follows from (a) and the other from (b). These conditions amount to an insightful analysis or decomposition of object dependence in experience.

The rest of this introduction gives background that's only briefly summarized in the dissertation and outlines my main contributions. It provides the big picture missing in chapters 2–5, which (as noted) are written as independent research papers. I've organized the introduction thematically, covering perception-based thought before moving onto implications for the nature of experience. First, §1.2 covers material related to question 2, on how perception allows you to select what's perceived for thought (chapter 3). Next, §1.3 covers material related to question 3, on experience's role (chapter 5). Last, §1.4 covers question 1 on object dependence. The first part, §1.4.1, covers the argument for object dependence (chapter 2). The second, §1.4.2, covers the extension to naïve realism (chapters 4 and 6).

1.2 Perception-based Thought

Conscious perception allows you to select what's perceived for thought. I first give three examples (§1.2.1). Then using an example I explain how conscious perception is a demonstrative means of selection (§1.2.2). Next, using another example I introduce relational means of selection and explain why you might think that demonstrative means of selection are relational (§1.2.3). This leads to the final topic, the perceptual link problem (§1.2.4).

1.2.1 Preliminary Examples

Consider the following three examples.

Example 1:
On a table in front of you sits a standard tennis ball. You look directly at it (you foveate the ball and hold your attention on it) and judge that it's green.

Example 2:
You hear a sudden and unexpected loud shattering sound in the next room. You wonder whether the noise was caused by your cat knocking your glass off the table.

Example 3:
With your eyes closed you feel around an object that's been handed to you. After feeling the texture, shape, and weight of the object, you judge that it's a drinking glass.

Call thoughts like these in which you exploit your conscious perception of a thing to select it for thought *perception-based thoughts*.[1] Let me make three points.

First, notice that in each case it would be natural to express the thought using a demonstrative term like 'that'. For example, in (1) you would naturally express your judgment by saying "*that* [pointing to the ball] is green". In (3) you would naturally express your judgment by saying "*this* [shaking the object] is a glass".

Second, all three examples involve propositional acts. They're episodic mental actions or mental events, not standing mental states like propositional attitudes. For example, (1) is an instance of judging and (2) is an instance of wondering. Other propositional acts include considering, entertaining, supposing, intending, inferring, and reasoning (engaging in a string of inferences). Propositional attitudes like beliefs and desires can also be perception-based. Presumably we form some beliefs and desires about things we consciously perceive by exploiting our conscious perception of them. More broadly, when I use the word 'thought' I have in mind both propositional acts and propositional attitudes.[2] Whether other kinds of thinking, e.g. imagistic imagining, can be perception-based is a question I set aside.

Third, each of the three examples has as its content a proposition with simple subject-predicate structure. For example, in (1) the property *green* is predicated of the ball at which you're looking. None of the examples, in contrast, involve contents with complex truth-functional structure (e.g., involving conjunction or negation) or quantification. While thoughts with contents having

---

[1]Smithies (2011b, 7) reports that the idea of perception-based thoughts, which I use here, goes back to Evans (1982). Both Evans and Smithies call them demonstrative thoughts.

[2]I'm using the term 'thought' to refer to states of, processes in, or events involving people. These states/processes/events have propositions as content and might involve the tokening of representational vehicles within the head (e.g., in a "language of thought", see Fodor 1975, 2008). While you could use 'thought' to refer to these contents or vehicles, I never have that use in mind.

complex structure or involving quantification can be perception-based (i.e., have perception-based components), for simplicity I shall only discuss perception-based thoughts with contents with simple subject-predicate structure. The subject of the thought—the thing to which a property is predicated—is the *target*.[3] What I mean when I say that perception-based thoughts are those in which you exploit your conscious perception of a thing to select it "for thought" is that you select it as the target of thought. For example, in (3) it's your conscious tactile perception of what you're holding which fixes it as the target to which you predicate the property of *being a drinking glass*.

1.2.2 Demonstrative Thought

Selecting a target by exploiting conscious perception is a demonstrative means of selection. Equivalently, perception-based thoughts are demonstrative thoughts. Calling perception-based thoughts demonstrative highlights an intuitive similarity, already noted above, between them and demonstrative speech acts. One suggestion locates this similarity in indexical representational vehicles: demonstrative speech acts involve demonstrative terms like 'that', and perception-based thoughts involve tokening analogous "mental demonstratives" in the head. A second suggestion locates it at the level of content, suggesting demonstrative speech acts and perception-based thoughts share the same kind of content (e.g., Peacocke 1983, ch 5). While both suggestions might be true, both are substantial theses which I won't assume here.

Before suggesting an alternative, it's useful to say more about mental demonstratives. Although I don't assume that perception-based thought involves tokening them, the view is common (e.g., Bach 2010, 55; Levine 2010, 179; Recanati 2012, 57–67). The view assumes a broader picture of thought on which *selecting* a target involves tokening a mental representation of that target. Mental representations are literal physical symbols in the head that play some role in mental activity and states, such as sensory processing, rational judgment, belief, memory, and motor control (Dretske 1988, 52; Cummins and Poirier 2004, 21). (Although mental representation *types* are usually defined psycho-functionally, their *tokens* are physical events or states in the brain.) A representation—not just a mental one—is *indexical* iff, for each token of that representation, its referent is determined by the context of tokening. Demonstrative terms in natural language are indexical in this sense. For example, a token of 'that' refers to whatever is

---

[3]I've adopted the term 'target' from Johnston (2011, 173) as it's a more convenient and less-loaded term than 'intentional object' or 'object of thought'.

being demonstrated (attended, pointed at, etc) by the speaker. On the mental demonstrative view of perception-based thought, you select what's consciously perceived for thought by tokening an indexical mental representation—called a *mental demonstrative*—the tokens of which have the functional role of referring to whatever is being consciously perceived.

Instead of identifying the similarity between perception-based thought and demonstrative speech acts in the kind of representational vehicles or content involved, it's more helpful—and less controversial—to identify it in the kind of capacity deployed. As Mike Martin points out (2002a, 178–81; see also Evans 1982; Recanati 2012, 62), both selecting a target for thought by exploiting conscious perception and referring in a speech act using a demonstrative term involve exercising a one-off, episodic capacity.

> **Example 4:**
> There are many everyday objects, properties, and locations with with you're acquainted, e.g. the car you drive, the town in which you live, and your friends and relatives. You think about these things all the time without exploiting your perception of them. (You often think about them when you're not currently perceiving them.) Right now, for example, I'm thinking about how my notebook is almost full.

When I have this thought I exercise a standing capacity to select my notebook. In contrast, when you select the tennis ball in example 1 you don't use any standing capacity. Instead, you select it in a way that depends on your current conscious perception of it.

Standing capacities to select targets are underlain by concepts. From my acquaintance with my notebook I have a concept, MY NOTEBOOK, and in example 4 I select the notebook for thought by deploying that concept. This concept picks out a specific thing—my notebook—and (like all conceptual thought) the target of the thought deploying the concept is the thing picked out. We need an account of concepts and what it is to deploy one in thought (for an introduction see Peacocke 2009; see also Peacocke 1992b; Fodor 1998; Prinz 2002; Machery 2009), but an intuitive understanding and the distinction between one-off and standing capacities suffices for explicating demonstrative thought. To summarize, selecting a target by exploiting conscious perception is a demonstrative means of selection in the sense that it involves a one-off, episodic capacity to select the perceived thing. It does not involve deploying a concept. Instead, conscious perception itself suffices for selection: you select what's perceived as a

target for thought by *directly* exploiting your conscious perception of it.[4]

1.2.3   Relational Means of Selection

A plausible suggestion is that perception-based thought is relational. This is a *metasemantic* claim about how the target of thought is fixed. It's helpful to begin by contrasting relational thought with descriptive thought.

> **Example 5:**
> Most likely you've never seen the first tennis ball that was produced and are not acquainted with it in any way. But you can think thoughts about it. You might wonder, for example, whether the first tennis ball produced was green and fuzzy.

Although you're not acquainted with the first tennis ball produced, you have the concept FIRST TENNIS BALL PRODUCED, and having a concept allows you to select that concept's referent. The concept FIRST TENNIS BALL PRODUCED is descriptive, i.e. it picks out a referent by describing it. Its referent is fixed by a condition it sets out, the condition *being the first tennis ball produced*. Whatever satisfies that condition—it happens to be the first tennis ball produced—is the referent of the concept and thereby the target of thoughts which deploy that concept.

In contrast, the conceptual thought in example 4 is not descriptive. My notebook isn't the referent of my concept MY NOTEBOOK because it satisfies some descriptive condition encoded in the concept. Instead, my acquaintance with the notebook, which lead to forming the concept, *relates* me (or the concept) to it. This relation of acquaintance, not any descriptive condition, fixes the referent of my concept MY NOTEBOOK, thereby fixing the target of my thought which deploys the concept. Instead of a descriptive concept, MY NOTEBOOK is an acquaintance-based concept.

Although not involving the deployment of acquaintance-based concepts, the perception-based thoughts in examples 1–3 are also relational. Take example 1. Your conscious perception relates you to the ball and it's that relation which fixes the ball as the target of thought. This suggests an explanation of what you

---

[4]This raises a final point. It's important not to confuse perception-based thought with thought that deploys perception or acquaintance-based (or "demonstrative") concepts to select a target. Often consciously perceiving something provides you with a concept of it, i.e. a standing capacity to select it for thought. Perception-based thought is something different from thought that exploits these acquaintance-based standing capacities, since it works in a way that depends on actually perceiving the target.

do when you (directly) exploit your conscious perception to select what's perceived as the target: you directly exploit *the relation* of conscious perceiving. A full picture of perception-based thought requires an account of what it is to *exploit* the conscious perception relation directly to select a target. One account is provided by mental demonstratives: exploiting the conscious perception relation involves tokening a mental demonstrative the functional role of which is to refer to whatever is being consciously perceived. But for now the point is just that in perception-based thought what makes the thing perceived the target is not that it satisfies some descriptive condition, but instead the way it's related to you.

The claim that examples 1–4 involve a relational means of selection distinct from the descriptive means in example 5 is controversial. Perhaps all concepts, even those formed via acquaintance, are descriptive. Maybe perception-based thought, although demonstrative and not mediated by concepts which encode them, still relies on descriptive conditions. Targets of perception-based thoughts might be fixed by the condition *the thing now consciously perceived*. I won't argue against these suggestions in this introduction. But note that that perception-based thought is demonstrative—that it involves one-off, episodic capacities to select—supports the claim that it's relational. This move from perception-based thought being demonstrative to it being relational is an important part of the argument in chapter 2 for object-dependent experience. Further, the move is intuitively plausible.

The idea is that the relational character of perception-based thought explains its demonstrative character. Although it might do it in other ways too (e.g., by acquainting you with the thing and thereby affording you an acquaintance-based concept of it, see fn. 4, page 8), conscious perception allows you to select what's perceived in a way that depends on your current perception. This is the sense in which it's demonstrative: it affords a one-off capacity to select the target. The best explanation for this dependence of selection on the relation of conscious perception is that the selection is relational. If exploiting conscious perception to select a target means exploiting the relation of conscious perception itself, then exploiting conscious perception will depend on your current perception of what's perceived. It's worth at least pointing out that if perception-based thought involves the tokening of mental demonstratives (again, something I don't assume), then it follows immediately that it's relational. This is because the referent of a token mental demonstrative, and thereby the target of thought, is fixed by "the context" of tokening, i.e. fixed by the relation of conscious perception.

1.2.4 Attention and the Perceptual-link Problem

Assume that perception-based thoughts are relational. There is a major issue lurking in the area. Above I assumed that the conscious perception relation is exploited. But if conscious perception is a relational means of selection, this can't be correct. Consider example 1. When you look at the ball on the table there are other things you consciously perceive at the same time. For example, you see not only the ball, but also the table on which it sits and things which are, from your perspective, behind the table. You might also, at the same time, hear the voice of someone next to you, tactilely perceive the floor under your feet, or smell the odor of some nearby food. So what makes the tennis ball the target of your thought can't just be that you're consciously perceiving it, since you're consciously perceiving many things. Conscious perception is a relation in which you stand to the tennis ball, but it's a relation in which you stand to many other things at the same time. So the relation you exploit in example 1 must involve something more than just conscious perception. The failure of the conscious perception relation itself to serve as the relation exploited in perception-based thought and the resulting need to specify the exploited relation I'll call the *perceptual link problem*.

A major thesis of this dissertation is that attention—specifically, voluntary focal attention—to what's consciously perceived is the exploited relation. Unlike conscious perception itself, it relates you to the target in a way that's exploitable. What matters in example 1 isn't just that you consciously perceive the ball, but that you consciously perceive the ball and voluntarily keep your attention focused on it.[5] Attending is a person-level action (not a variety of subpersonal cognitive processing) directed at perceived objects, properties, and locations. It's something you do. Three examples that involve attention are visually tracking an object as it moves past you, visually searching a scene for a specified target, and being startled by a loud bang or bright flash. In the first case you hold attention on the moving object, in the second you shift attention around the scene, and in the last case attention is grabbed by the sudden noise or flash. The first two are cases of voluntary attention, the last involuntary. If during the visual search you carefully attended to each object in the scene one at a time your voluntary attention would further be focal, or focused. (Alternatively, you might scan multiple

---

[5]You might suggest that although you stand in the conscious perception relation to many things, what you exploit is the one *instance* of that relation between you and the tennis ball. Perhaps all attention does, the suggestion goes, is fix which instance of the conscious perception relation is exploited. But there doesn't seem to be any substantial difference between this suggestion and my proposal.

objects at once, leaving attention spread.)

The main competitor to my proposal is the dominant causal source view (see Evans 1973, 1982; Raftopoulos and Müller 2006, 253).[6] On this view perception-based thought involves the tokening of a mental demonstrative which is, or is a label for, a mental file (Bach 1987; Raftopoulos and Müller 2006, 256–8,264; Dickie 2010, 222, 2011, 305; Jeshion 2010c, 135; Recanati 2010, 157, 2012, 34–8, 57–67). A *mental file* is functionally defined as a collection of information or beliefs which our cognitive processes treat as being about the same thing (Bach 1987; Raftopoulos and Müller 2006, 256–58,264; Pylyshyn 2007, 37; Dickie 2010, 222–25, 2011, 304; Jeshion 2010c, 129–35; Recanati 2010, 156, 2012, 34–38). A mental file refers to what it's about, and it's about the dominant causal source of information in the file.[7] On one possible view there is a special-purpose temporary file that's opened to store incoming perceptual information and closed (and reopened) as what's being perceived changes. This file is the mental demonstrative tokened (by closing and opening the file) in perception-based thought. Being the dominant causal source of information contained in it is a way of relating to this special-purpose file/perception-based mental demonstrative. This relation, according to the dominant causal source view, is what's exploited in perception-based thought.

The perceptual link problem and the proposal that voluntary focal attention is the relation exploited in perception-based thought are more fully explicated in chapter 3. There I explain voluntary focal attention in more detail. I then argue that my attention-based proposal better fits the intuitive data than the dominant causal source view. There are cases of perception-based thought in which what's focally attended come apart from the dominant casual source and, intuitively, what's focally attended is the target of thought. Finally, I argue that the restriction to voluntary attention is necessary, since involuntary attention is often divided between multiple objects even when, intuitively, we can use voluntary attention to select something consciously perceived.

---

[6]Most statements of the dominant causal source view are ambiguous between it and an attention-based view like the one I defend. One contribution of chapter 3 is to clearly distinguish the views.

[7]This part of the view is optional. For example, you could combine a mental file account with my attention-based attention and posit that the relevant mental files involved in perception-based thought are about whatever you're currently focally and voluntarily attending.

1.3 The Role of Experience

Question 3 from §1.1 asks what, if any, role is played by experience in perception-based thought. After covering some preliminaries (§1.3.1), I explain the motivation for thinking experience has a role at all and then give my account of that role (§1.3.2). I conclude by discussing how my work goes beyond the accounts from Campbell and Smithies already in the literature (§1.3.3).

1.3.1 Experience, Content, and A-Consciousness

Particular episodes of perceiving are events involving transmission-transduction-neural interactions between you and what's perceived. Take the example of perceiving the tennis ball (example 1). Light reflects from the ball to your eyes, its energy is transduced into electrical signals, and information about the ball is processed in "streams" of connected neurons. The processing might—as the dominant view in cognitive science holds—involve the construction of a representation of the ball, but this is controversial (see Marr 1982; Pylyshyn 1986, 2007, 1–9; responses include Gibson 1966, 1986; Brooks 1991; van Gelder 1995; Noë 2004; Hutto and Myin 2013; Orlandi 2011a,b, 2014; see Eckardt 2012 for overview). The interaction possibly involves larger body movements, e.g. you repositioning your head or body to get a better view of the ball. (I take this approach of explicitly thinking of perceiving in these terms from Johnston 2011, 172–77, but the idea that perceiving involves, more broadly, a certain kind of causal interaction with the world is widespread.)

In typical cases perceiving involves an experience or, alternatively put, has phenomenology. I give two broad ways to characterize experience (both from Nagel 1974, 437–38), then ostensively point to it with examples. (See Siewert 2012 for a more detailed characterization.) First, the experience is the first-person perspective or subjective aspect of perceiving. Second, there's "something it's like" when you perceive, or, alternative put, there's some way what's perceived looks to you visually, or sounds to you auditorily, etc. The experience is the particular *way* an episode of perceiving "is like", subjectively.

Next, the following examples highlight experience by presenting contrasting cases between which the experience changes.

Example 6:
Assuming you wear glasses, look at an object in front of you and contrast how it looks with your glasses and without your classes. The way it looks to you changes when you take your glasses off—it gets

blurry. This is a change in experience.

Example 7:
Look at a nearby object. Changes in lighting conditions and the angle at which you view it will change how it looks to you, in some sense. For example, a penny straight on looks circular, but looks elliptical from most perspectives. These are changes in experience.

The above two examples both involve subtle differences in experience. But it's also helpful to point to big differences.

Example 8:
It seems conceptually possible that the way some people see colors is inverted. Perhaps, unknown to us, the way red things look to you is the way green things look to me. If so, this would be a (radical) difference in experience.

Example 9:
In example 3 you feel the roundness of the glass. If you opened your eyes you'd see the roundness too. In both cases you perceive the roundness of the glass, but the experience is very different. What it's like to feel roundness is different from what it's like to see roundness.

As the next example shows, just perceiving different things usually leads to big differences in experience.

Example 10:
Contrast what it's like to see a tennis ball on a table in front of you with what it's like to see a tree. Obviously it's quite different. Tennis balls (normally) look quite different from trees.

The final example involves not differences in experience but sameness:

Example 11:
Compare what it's like to see an apple with what it's like to see a toy truck of the same color. The experience in each case is different overall, but still similar in one obvious way. They both look red. In that respect the two episodes of perceiving are the same experientially.

Note that the initial characterizations were used in the examples to help focus on the experiential difference: the relevant differences are always differences from your perspective, in the subjective aspect, or in "what it's like".

An important point is that experiences are things which particular episodes of perceiving instance; they are a way of typing episodes of perceiving (e.g., Siegel 2010a, 20).[8] Let me expand on this point while running through some related jargon. You have a certain experience when you're the subject of an episode of perceiving which instances it. Something is consciously perceived if you have an experience while perceiving it. Two episodes of perceiving involve having the same experience, or have the same phenomenology, iff they are the same subjectively. There's a change in experience, or in phenomenology, between two episodes of perceiving when the two differ subjectively, i.e. when they instance different experiences. For example, there's a change in your experience before and after you take off your glasses in example 6. While I'm talking as if token experiences—particular instances of some experience—*are* episodes of perceiving (Byrne 2009, 431–35; Johnston 2011, 180), nothing below depends on this identification.[9]

Along with having experience, perceiving involves content.[10] Perception always involves perceiving the world as being a certain way. The way the world is perceived as being in an episode of perceiving is its *content*. For example, when you perceive the tennis ball in example 10 it looks as if there's a round, green, fuzzy object of some size a certain short distance in front of you. Perceptual content might not be accurate. For example, if the lighting is just right the tennis ball may actually look blue (instead of merely looking like a green ball under off-white light). When perceptual content doesn't match up with the world the perceiving is illusory, or an *illusion*. When it does the perceiving is *veridical*.

At least as a conceptual point, experience and content are dissociable. Differences in experience and content often go together, e.g. in example 10 both the

---

[8]This way of distinguishing between the (by definition) world-involving episodes of perceiving and the experiences they instance is now common (e.g., Siegel and Silins forthcoming).

[9]On the distinction between experience (i.e., experience types) and token experiences, see Pautz (2009, 494), Burge (2005, 34), and Martin (2002a).

[10]It's widely held that perceiving involves content (for early statements of the view, see Evans 1982, 226; McGinn 1989, 58; Harman 1990, 43; Davies 1992, 22; Peacocke 1992b, 61; McDowell 1994b, 11,26; for general discussion see Nanay forthcoming). But the view is not universally accepted (e.g., Campbell 2002; Travis 2004; Fish 2009; Brewer 2011). The view that perception, or perceptual experiences, have content is related to, but need not be identified with, the representational view of experience. See §1.4.2, §2.7, especially fn 10 (page 46), and §6.2 for discussion.

experience and content of perceiving a tennis ball differ from those of perceiving a tree. But—as a plausible suggestion—there can be changes in experience without changes in content. For example, when you take off your glasses in example 6 there is a change in experience, but not obviously a change in content. What you're looking at may now look blurry, but it's not as if you perceive the world to be different. Similarly, as you change your viewing angle in example 7 the experience changes (e.g., the way a penny looks when you view it straight on is different from how it looks, in some sense, when you view it at an angle). But, plausibly, merely shifting viewing angles doesn't change perceptual content: when I move it doesn't look to me as if what I'm perceiving changes its shape. Further—again as a plausible suggestion—there can be changes in content without changes in experience.

Example 12:

You look at a #2 pencil and then with your eyes closed it's swapped with another identical pencil. You open your eyes and look at the new #2 pencil.

In this example the content when you perceive the first pencil differs from the content when you perceive the second for the simple reason that you're perceiving different things.[11] But, since the two are identical everything will look the same to you; the experience won't have changed.

Often perceptual content is attributed to the experience: it's the experience enjoyed in an episode of perceiving which is illusory or veridical. Perceptual experience, it's said, "has" the content. But this builds in that there can't be changes in content without changes in experience. You might suggest attributing content to token experiences instead, but this (as does the first suggestion) suggests that only episodes of perceiving which involve having an experience have content. Since, as we'll see, it makes sense to talk about the content of perceptual episodes which don't (see §5.2), we might attribute the content to the episode of perceiving itself. But this has a disadvantage as well, since when there is an

---

[11]This example won't work if you deny that the content of experiences involves the particular objects perceived (e.g., Davies 1991; Tye 1995; Pautz 2009), but most philosophers now accept that it does involve particular objects. Unfortunately simple, intuitively convincing examples which don't involve a change in the particular object perceived are hard to find. In chapter 2 I'll argue that this example actually involves a change in experience, but here it suffices to demonstrative the conceptual coherence of changes in content without changes in experience. See §1.4.1 for an overview.

experience there must be some connection between it and the content. For example, if you see the tennis ball in example 10 and have the same experience you do when you see the tree, then it's unintuitive to attribute veridical content to the perceptual episode, e.g. to say you perceived the world as containing the tennis ball. Whether perceptual episodes, token experiences, or experiences primarily have content won't matter below, but that there's an issue here is worth flagging.

When perceptual content is available for spontaneous, deliberate use in other tasks (e.g., answering questions or navigation), the episode of perceiving is *access conscious* (see Block 1995). Alternatively, we might say the perceptual content is access conscious. At least as a conceptual point experience and access consciousness are also dissociable. For example, inattentional blindness might provide cases in which an episode of perceiving involves having an experience without being access conscious (for work on inattentional blindness, see Rensink et al. 1997, 2000; Mack and Rock 1998; see Jensen et al. 2011 for review). In inattentional blindness, while focusing your attention on one object others in the same spot or nearby go unnoticed. It seems plausible that the unnoticed objects are experienced—there's some way they look to you, or some subjective point of view you have on them—but your failure to notice them suggests the perceptual content was not access conscious. Also, an episode of perceiving can be access conscious without involving an experience. Blindsight, discussed next, provides an example.

1.3.2 The Attention-Guiding Role

Why think that experience has a role to play in perception-based thought at all? The main reason is that it seems to be necessary: while consciously perceiving something makes it available for thought, it's intuitive that nonconsciously perceiving it does not. You would nonconsciously perceive something if you perceived it, including having perceptual processing of sensory input from it, but (for whatever reason) there was no experience of what's perceived. Blindsight, an actual clinical condition, provides a concrete example (see Brogaard 2011a; Brogaard 2012 for discussion).

In humans and other primates about 90% of the ganglion cells in the central visual pathway project from the retina to V1 (the primary visual cortex) through the dorsal lateral geniculate nucleus (Weiskrantz 2009, 69). In humans the destruction of some portion of V1 leads to total blindness (the absence of visual experience) in the part of the visual field to which that portion of V1 maps topographically. But in cases in which the damage is restricted to V1 (leaving the

areas of visual processing beyond that intact) some residual visual functioning often still remains, leading to a loose collection of abilities called *blindsight*. For example, in some experiments D.B. (one of the original people in whom blindsight abilities were discovered) was shown a flash in the blind patch of his visual field and was asked to either direct his gaze or point to the location he *guessed* the flash had been. This he could do well above chance (Weiskrantz 2009, 87–93).

The standard interpretation of blindsight is that the residual visual functioning is explained by the small portion of the optic nerve (∼10%) that lies outside the central visual pathway. Ganglion cells from the retina outside this pathway project to extrastriate cortex (V2–V5), inferior temporal (IT) cortex, and subcortical regions via 9 or 10 other pathways (Weiskrantz 2009, 59,69). The idea is that V1 is necessary for conscious visual perception, but these additional pathways facilitate some level of visual processing that still allows for some limited forms of functioning. While there's no conscious visual perception of objects within the blindspot, there is still some visual states which represent those objects and have some influence over behavior (Block 1995, 230).

Now imagine that you're a blindsighter. You are, from your point of view, blind in some area of your visual field; you have no visual experience of things which fall within it. But nevertheless you still visually perceive what's there, in the sense that there's processing of visual sensory stimuli in that area. This processing or nonconscious perception even affords limited functioning. Like D.B., if you're asked to guess and point to the location of an object within that blind spot you can do this with near perfect accuracy. When forced to guess about other features of the object like its shape and color you are nearly always correct. Of course, from your point of view you're making guesses about something you can't see. You have no visual experience of the object—there's just nothing there.

The intuitive suggestion is that in this case you could not exploit your nonconscious perception of the things in your blindspot as a means to select them for thought (Campbell 2002; Johnston 2006; Smithies 2011b). You may still be perceiving what's there, and even have the guessing ability described, but without actually having visual experiences of what's there—without consciously seeing them—you can't exploit that perception. Instead, it seems that if you wanted to select what's in your blindspot for thought, then you would have to use a descriptive concept. For example, you may use the concept THING I'M PERCEIVING IN MY BLINDSPOT to select it. Similarly, the suggestion goes, no case of nonconscious perception can be exploited to select what's perceived as the target of thought.

Two points should be noted. First, the claim that blindsighters cannot exploit their perception of what's in their blindspot is based on intuition (albeit an intuition widely shared among those who work on the topic). It's open for someone to imagine themselves in the position of the blindsighter and to conclude that she can exploit her nonconscious perception.

Second, blindsighters not only lack experience of the things they perceive in their blindspot, but they also lack access consciousness to the contents of that perception (see Block 1995). While the perceptual content does influence the blindsighter's behavior, it isn't available for spontanous or deliberate use. The blindsighter only succeeds when she is given prespecified choices between which to guess. Even if you agree with the intuition that blindsighters cannot exploit their perception of what's in their blindspot, you might suggest that the problem isn't their lack of experience. Perhaps, the suggestion goes, nonconscious perception in which you're still access conscious of the content would allow for selecting what you unconsciously perceive (see Kelly 2004, 283–84 for this suggestion). If so, then the initial motivation for thinking that experience has a role in perception-based thought—that it's necessary—is undercut.

In chapter 5 I argue, focusing on the case of vision, that experience has the role of making perceptual content available for use in voluntarily directing focal attention. That is, experience has the role of making the perception's content access conscious—at least access conscious for use in the voluntary control of focal attention. Because my argument does not depend on intuitive considerations about blindsighters I avoid the two issues just noted. Instead, I offer two lines of evidence. The first is introspective: I claim that when we introspect what it's like to voluntarily guide attention it's plausible that we rely on our experience. The second is neurobiological. I argue that the way in which visual processing areas most closely related to conscious visual perception (the inferior temporal cortex) project into the area of the brain responsible for voluntary control of attention (the frontal eye fields, FEF) support the idea that conscious perception in particular has a role to play in making perceptual content available for use in the voluntary control of attention.

1.3.3   Other Accounts

Both Campbell (1997; 2002; 2004; 2011b) and Smithies (2011b) have also provided accounts of experience's role in perception-based thought.[12] On Smithies'

---

[12]Smithies very clearly has in mind perception-based thought. But Campbell oscillates between talking about three distinct things: (1) experience's role in perception-based thought,

account, experience plays an epistemic role. Specifically, "it enables subjects to use demonstrative concepts in forming immediately justified beliefs about objects in the world around them" (2011b, 19). I'm going to set Smithies' account aside for two reasons. First, even without details it's clear how the account differs from mine. On my account experience has what Smithies calls a target-setting role. Voluntary focal attention sets the targets of perception-based thought (by providing an exploitable relation), and experience has the role, at least in vision, of aiding this target setting by making perceptual information available for use in the voluntary control of attention. Second, because the roles are so radically different the two accounts are not necessarily competing. Experience could play both an epistemic role and a target-setting role in perception-based thought.

This leaves Campbell's account which, as Smithies notes, also gives experience a target-setting role. My aim here is only to summarize the relevant details of the account and how it differs from mine. First, Campbell's account assumes that your knowledge of the reference of demonstrative terms, and presumably your selection of targets in perception-based thought, "is what causes and justifies the use of particular procedures to verify and find the implications of propositions containing the demonstrative" (Campbell 2002, 25). So, whatever provides this knowledge, or allows for selection of targets in perception-based thought, must be the sort of thing which can fill this functional role of causing and justifying the use of verification procedures (2002, 26; see also Smithies 2011b, 14). Campbell claims that what provides this knowledge, or allows for selection, is conscious attention to the target (2002, 25, 2004, 268). By 'conscious attention' Campbell seems to mean voluntary attention to consciously perceived (experienced) objects.

Two points should be noted. First, so far Campbell's account looks like mine.

---

(2) experience's role in providing perception or acquaintance-based concepts (see fn. 4, page 8), and (3) experience's role in providing knowledge of the referents of demonstrative terms. Here are examples, in reverse order. Campbell says, "you make a remark to me about 'that woman' ... It is only when ... I focus on that person, that I would ordinarily be said to know who was being referred to" (2002, 8–9). He also says, "we cannot understand how experience, so conceived, could be what provides us with our concepts of the objects around us" (2002, 123). Finally, he says "experience of objects has to explain our ability to think about those very objects" (2002, 114). Sometimes Campbell shifts between these different things from one sentence to the next. Textual evidence also suggests that when he talks about providing concepts, he really does mean providing acquaintance-based concepts (see, e.g., 2002, 138). I merely flag the issue here and set it aside. Some interpret Campbell as having perception-based thought in mind (e.g., Clark 2006; Smithies 2011b), while others take him to have in mind experience's role in providing perception or acquaintance-based concepts (e.g., Campbell 2011a; Cassam 2011; Roessler 2011).

I also claim that it's voluntary attention to experienced objects which sets the targets of perception-based thought. Second, because Campbell makes this assumption about the functional role of whatever sets the targets in perception-based thought, he incurs the explanatory burden of explaining how voluntary attention to experienced objects causes and justifies verification procedures. To do so, he gives a detailed account, based on Treisman's work on attention and feature binding (Treisman and Gelade 1980; Kahneman et al. 1992; Treisman 1998), of how this happens. One way my account differs from Campbell's account is that I don't accept the assumption about the target-setting role and therefore don't try to give any account of how voluntary attention to experienced objects might cause and justify verification procedures.

A third point to note that, at least in his earlier work (2002; 2004), Campbell specifies that voluntary attention to experienced objects plays the target setting role, but isn't careful to specify the specific role of experience itself. For example, I claim that experience has the specific role of making information available for use in voluntarily guiding attention. Where Campbell does focus on the role of experience itself (2002, ch 7), he focuses on its role in providing perception or acquaintance-based concepts. So, a second way in which my account differs from Campbell's account is that I further specify the role of experience itself within the broader target-setting role of voluntary attention.

In later work (2011b) Campbell does clearly specify the role of experience itself within attention. Here his account sounds similar to mine. He says that it's consciously experienced properties which are used to select targets of attention. This leads to the third substantial difference between my account and Campbell's account. Even if Campbell and I agree on the role of experience itself, my argument for this role is different and original. My introspective (§5.3.2) and neurobiological (§5.3.3) arguments are different from any of Campbell's arguments.

A fourth difference between my account and Campbell's account, which I discuss at some length in chapter 5, is that he holds that only experience can play the role it actually does in perception-based thought. In contrast, I argue that there's evidence that nonconscious perception could also play the role of making information available for use in the voluntary control of attention. The neurobiological evidence I discuss is key to making this point.

1.4 Implications

1.4.1 Object-dependent Phenomenology

Example 12 suggests that if you have two visually identical #2 pencils—you cannot discriminate them by how they look to you—then two episodes of perceiving which differ only in which pencil you see will instance the same experience. Holding lighting, surrounding background, and your point of view constant, swapping pencils will not change your experience. This, anyway, is the standard thing to say about the example. Usually it's supported by the following reasoning (e.g., Chalmers 2006, 108; Siegel 2010a, 169; Schellenberg 2011, 738; see Martin 2002a for overview):

1. Since the two pencils are visually indiscriminable, if the only change is a swap in the pencils then the two token experiences (the one before, and the one after the swap) are also indiscriminable.
2. If two token experiences instance different experiences (i.e., differ phenomenally), then they're discriminable.
3. So, the two token experiences instance the same experience (i.e., do not differ phenomenally).

Johnston (2011, 181) calls (2) the phenomenal bottleneck principle. If it's true, then this reasoning shows that experiences are *object independent*. That is, there can be changes in what's perceived without a change in experience. Equivalently, two tokens of an experience can involve perceiving different objects.

I argue against this picture in chapter 2.[13] Experiences are *object dependent*: if there's a change in the object perceived, then there's a change in experience. The argument (as noted in §1.1) depends on the assumption that perception-based thoughts are demonstrative and that experience plays a role in them.

> **Thesis of Chapter 2**
> If (a) perception-based thought is demonstrative, and (b) experience plays some role in it, then experiences are object dependent.

Chapters 3–5 argue for a particular account of perception-based thought on which (a) and (b) are true. But, as §1.2 and §1.3 show, both (a) and (b) are intuitively plausible and widely held. So the argument in chapter 2 forces a dilemma if you deny that experiences are object dependent: you must also either deny that

---

[13]There are other approaches to arguing against the object dependence of experience (see Lycan 2001; Fisher 2007; Mehta 2013).

perception-based thought is demonstrative, or that experience plays some role in it.

Although they do not articulate or argue for the thesis of chapter 2, McDowell (1986) and Campbell (2002, 125; 2010, 201) advance the following similar thesis (see also Brewer 2006, 179; Johnston 2006, 264; compare with Pautz's similar "grounding principle" 2009, 500).

Basic Object-dependency Thesis
Conscious perception introduces what's perceived for thought only if experiences are object dependent.

They argue that were experiences not object dependent, then conscious perception would lack something required for it to introduce the thing perceived for thought. It would lack "directedness towards external objects" (McDowell 1986, 165) or would fail to "present" determinately a particular object (Campbell 2002, 125). As Martin noncommittally explains while elucidating McDowell and Campbell's insight (2002a, 198), perceiving the second #2 pencil obviously doesn't introduce the first for thought. But if your experience when perceiving the first #2 pencil is the *same* as your experience when perceiving the second, then, surely, conscious perception doesn't introduce the first even when that's the one actually perceived.

The above merely gestures at an intuitive puzzle for explaining how conscious perception introduces what's perceived for thought if experience is object independent. It's the one with which this introduction opened (see question 1, §1.1). Although they give some additional details (which I won't try to reconstruct), McDowell and Campbell do not go beyond this intuitive puzzle (see Burge 2005, 60,62; Pautz 2010, 283,286). The problem with leaving it there is that there's a plausible account of how conscious perception introduces what's perceived for thought—the one from §1.2—which, prima facie, doesn't require the object dependence of experience. So maybe the intuitive puzzle solves itself: once we get clear on how conscious perception introduces what's perceived for thought we'll understand why experiences need not be object dependent. This response comes from Burge (2005, 44,54–63; see also Martin 2002a, 178–81,197).[14] The thesis of chapter 2 and the argument for it largely fall out of seeing how the response fails.

---

[14]Martin (2002a) advances the discussion between Burge, McDowell, and (by extension) Campbell by setting out the issues more clearly than they had been before. My overview here of the basic object-dependency thesis owes much to this paper from Martin.

Before presenting Burge's response in more detail, note that McDowell and Campbell leave the antecedent phenomena in the basic object-dependency thesis underspecified. Conscious perception introduces objects for thought in two ways: by providing an exploitable means of selecting what's perceived as the target of a thought, and by providing concepts of perceived objects and properties (see fn. 12, page 19, and §1.2.2, especially fn. 4, page 8). Which is meant matters, since prima facie the requirements on experience for each might differ. Or, the reasons for those requirements might differ. My project advances the discussion by clearly separating these ways and focusing on one. It matters for the argument in chapter 2 that what's at issue is how conscious perception provides a means of selecting targets. The argument doesn't go through—or even make sense—if applied to how it provides concepts.

Returning to Burge's response, the antecedent of the basic object-dependency thesis says that conscious perception provides an exploitable means of selecting what's perceived as the target of thought. On the account in §1.2 conscious perception is a demonstrative means of selecting. Assume (as Burge does, 2005, 44) this means that conscious perception allows you to select what's perceived by relating you to it—by providing a perceptual link—in a way you can directly exploit for selection. Now the crux: why think the perceptual link is exploitable only if experiences are object dependent? For example, on the mental demonstrative account exploiting the link just involves tokening a mental demonstrative with the functional role of referring to whatever is linked. There's no obvious reason why this requires that experiences depend on what's perceived. Perhaps the exploitability of perceptual links only requires the object-dependence of the *content* of the associated token experiences (Speaks 2009, 560; Siegel 2010a, 157; Schellenberg 2011, 739). Or, perhaps the involvement of causal relations in the perceptual links is sufficient for their exploitability (Pautz 2009, 499, 2010, 286).

The insight behind the thesis of chapter 2 is that if the exploitability of perceptual links doesn't require that experiences are object dependent, then it's unclear whether experience plays a role at all. The worry is that if selecting a target by exploiting your perceptual link to it doesn't require that your experience depend on what's perceived, then nonconscious perception should be exploitable too. If the worry is right, then if perception-based thought is relational and nonconscious perception isn't exploitable (i.e., experience plays a role), then experience is object dependent. By further arguing that demonstrative thought is relational we get the thesis of chapter 2: (a) perception-based thought is demonstrative and (b) experience has a role to play only if experiences are object dependent.

The work of chapter 2 is to turn the worry about the role of experience into an argument. This brings us to another insight: we get the needed argument by decomposing object dependency into two jointly sufficient conditions about what's presented in experiences. I give these two conditions and then argue that the first follows from (a) while the second follows from (b). Experiences are object dependent if:

1. Presentation is object dependent: Necessarily, an experience can be instanced only when the object presented in it is actually perceived, and
2. Presentation is phenomenally relevant: Necessarily, all instances of an experience present the same object.[15]

That the conditions jointly suffice for object dependence is straightforward. Recall that experiences are object dependent if a change in the object perceived is sufficient for a change in experience (in phenomenology). By condition 1, if two token experiences differ in what's actually perceived they will present different objects. By 2, they don't instance the same experience; there's a change in experience.[16]

It takes substantial arguing to show that the object-dependence and phenomenal relevance of presentation follow, respectively, from (a) the demonstrative character of, and (b) experience having a role in, perception-based thought. I leave those arguments for chapter 2. But it's plausible at the start that the phenomenal relevance of presentation follows from (b). If two instances of the same experience can present different objects, then it seems as if the experience plays no role in how the instances provide a means of selecting what's perceived as a target of thought. Similarly, once it's established that demonstrative thoughts are relational, it's plausible that the object dependence of presentation follows from

---

[15]Applying the terms 'object dependent' and 'phenomenally relevant' to presentation is merely a convenient way to refer back to the two respective conditions without repeating them when something more descriptive than 'condition 1' and 'condition 2' is helpful.

[16]It should be clear that condition 1 isn't sufficient for object dependence, since without 2 it's left open that two token experiences might present different objects but still instance the same experience. But condition 2 isn't sufficient either. Consider the normal case of having a token experience of the first #2 pencil (one that presents it) while perceiving it. Case 1: you have an hallucinatory experience of the first #2 pencil when neither is there. Case 2: you have a token experience of the first #2 pencil while actually perceiving the second. Condition 2 alone leaves room for cases 1 and 2. But both would provide a pair of token experiences (the normal case vs case 1, and the normal case vs case 2) which differed in what's perceived, but which presented the same object (the first #2 pencil). So condition 2 would leave open pairs of cases which instanced the same experience, despite a difference in what's perceived. We need condition 1 to rule these out.

(a). If experiences can present an object without that object being perceived, then it seems as if their instances would not provide an exploitable relation to what's perceived and hence not enable demonstrative selection.

Whether they are object dependent is a fundamental question about experiences. First, if experiences are object dependent, then the phenomenal bottleneck principle is wrong. Two experiences can differ without being discriminable by the subject. Second, if experiences are object dependent then phenomenal internalism should be rejected in favor of phenomenal externalism. *Phenomenal internalism* is the view that the phenomenology of a particular instance of perceiving is fully determined by what goes on within the bounds of the biological systems of the perceiver, i.e. "within the head" (Burge 2005, 22, e.g., endorses the view). The particular object being perceived and the transmission of energy via reflected light to the eyes is only *causally* relevant to the experience enjoyed by the perceiver. The experience, on this view, is constitutively determined by the neural states and processing that happens after the transduction of energy by receptors in the eyes. In contrast, *phenomenal externalism* says that what goes on outside the biological bounds of the perceiver can constitutively (noncausally) affect phenomenology. There can be differences in experience even if there's no difference in the neural states and processing within the head. (See Gertler 2012 for some problems with this distinction.)

1.4.2 Naïve Realism

Views on the nature of perceptual experience divide into two groups: naïve realist and representationalist (see Genone forthcoming for discussion). Naïve realism takes experiences to be relations to what's presented in them (Campbell 2002, 2010; Brewer 2006, 2007, 2011; Johnston 2006, 2011; Fish 2009; Leddington 2009; Kennedy 2011, 2013; Logue 2012b,a; Hobson 2013; Knight 2013; Genone 2014). Representationalism takes experiences to be representations of what's presented in them (Burge 1991, 2005, 2009, 2010; Davies 1991, 1997; Tye 1995; Dretske 2003; Chalmers 2004, 2006; Byrne 2009; Pautz 2009, 2010; Speaks 2009; Schellenberg 2010, 2011, 2013; Siegel 2010a,b). Note that the claim that experiences are representational is equivalent to the claim that experiences have content. Representationalism is distinct from the weaker view that token experiences are representations. Token experiences might be representations, i.e. have content, even if experiences are not themselves representations. This would happen, for example, if there could be changes in experience without changes in content. Blurriness (example 6) has already been mentioned as a putative ex-

ample (see also Peacocke 1983, ch 1; Block 2010). Although naïve realists tend to deny this weaker token version of representationalism for other reasons (Campbell 2002; Travis 2004; Brewer 2006, 2011; Fish 2009; see fn. 13, 121, for details), there isn't even a prima facie conflict between their view and it. Experiences could be relations while their tokens are representations because their tokens might also instance types which are representations.

Arguments from experience's role in perception-based thought, like the one just discussed in the previous section (§1.4.1), have been put forward as supporting naïve realism (e.g., Campbell 2002, ch 6). But as I discuss in §2.7, the object dependence of experiences isn't sufficient for naïve realism. There are versions of representationalism which also make experiences object dependent. So, the argument from chapter 2 doesn't establish naïve realism. Further, there's no way to extend the argument to rule out the versions of representationalism on which experiences are object dependent. In this final section of the introduction I provide some preliminary background to the representationalism/naïve realism dispute, describe the versions of representationalism on which experiences are object dependent, then briefly summarize the results from chapter 4 and how chapter 6 extends them to rule out these versions of representationalism.

Naïve realism and representationalism are direct-realist reactions against the sense-data view (Siegel 2006b, 356, 2010a, 176; Logue 2009, 18; Brewer 2011, 11; see §6.1). Direct realism is the view that what's presented in token experiences are mind-independent objects. Specifically, what's presented is whatever is perceived. In contrast, the sense-data view says that token experiences present mind-dependent "sense data" and that you only perceive the mind-independent world indirectly through them (Russell 1997/1912; Price 1932; Ayer 1956; Jackson 1977). On the sense-data view, experiences are relations to what's presented in their tokens (sense data). Naïve realists keep the relational aspect of the view, but say that what's presented is the mind-independent world. Representationalism agrees with the direct realist aspect of naïve realism, but instead of relations says that experiences are representations of what's presented in them.

If you're a direct realist, the motivation for representationalism is that it allows for experiences to be object independent. As the argument from hallucination shows (see §6.1), indirect realism follows if experiences are both object independent and relational.[17] One way to do this is to make the content of expe-

---

[17]In §6.1 I show how both representationalism and naïve realism can be motivated as responses to the problem of hallucination. The considerations I present here about object independence are the representationalist half of that picture. The object independence of experiences is

riences existential (e.g., Davies 1991; Tye 1995; Pautz 2009). Consider example 12. According to existential-content representationalism, when you see the first #2 pencil your experience does not represent the pencil itself. Instead, it represents that there exists a yellow, thin cylindrical object 10cm long about a meter from you. Since the second pencil is identical to the first, your next experience has the same existential content. So they're the same experience.

Representationalism allows for object independence even if the content is singular (e.g., Soteriou 2000; Burge 2005, 2010; Speaks 2009; Schellenberg 2011). First, you could treat the singular element of experiences as an indexical representation. It's not that with each experience there's a particular object which it represents, but each experience has the functional role to represent what's perceived. So there's no change in experience when the pencils are swapped, but the token of that experience before the swap represents the first pencil and the token after the swap represents the second pencil.[18] Second, you could treat the singular element as nonindexical, but still detached from what's perceived. Each experience represents some particular object, but it might not represent what's perceived. Perhaps there's no change in experience when the pencils are swapped because your experience continues to represent the first even after the swap.

Although representationalism is motivated by the intuition that experiences are object independent, the view itself also allows that experiences are object dependent (see, e.g., Speaks 2009). This would happen if you opted for a singular content version of representationalism on which the content was both nonindexical and dependent on the object perceived (for object dependent representations, see Burge 1977, 346; Evans 1982; McDowell 1982, 204, 1984, 287; Martin 2002a, 178; Crane 2011, 23). On this view the experience itself, the repeatable type instanced in subjectively identical episodes of perceiving, represents the particular object presented in it. But it's not an indexical representation. It doesn't have context-sensitive content whereby instances of it represent what's perceived in them. Instead the experience itself representations some particular object. Further, the representation is dependent on that object. The experience cannot be instanced in a perceptual episode in which that object isn't perceived. Of course, this version of representationalism isn't consistent with the motivation for repre-

---

one of the two main premises of the problem of hallucination, the one which representationalists want to save.

[18]Note that the nonsingular elements of the experience aren't indexical. For example, if the first pencil was instead swapped with a green one the experience would change. The experience before the swap would represent that that (the first pencil) is yellow, while the experience after the swap would represent that that (the second pencil) is green.

sentationalism—which is to have a version of direct realism on which experiences are object independent—but the point is that it's not ruled out by the minimal representationalist commitment that experiences are representations.

Since chapter 2 shows that experiences are object dependent and object dependence rules out indirect realist views such as sense-data theory, establishing naïve realism only requires ruling out the object dependent version of representationalism. Chapter 6 rules out this version of representationalism by arguing that representationalist views themselves are committed to more than the basic claim that experiences are representations. Specifically, I argue that representationalists are committed to grounding the content of experiences in the neurologically realized representations constructed by our sensory systems. Chapter 4 in turn argues that you are presented with things in experience for which no neurologically realized perceptual representations are constructed. This makes it impossible that the content of experiences could be grounded in the way required by representationalism, thereby ruling out representationalism.

Chapter 2:
An Argument for the Object-Dependence of Experience from Demonstrative Thought

2.1 Introduction

The claim that experiences are object dependent is widely rejected. Instances of perceiving numerically distinct objects need not involve having different experiences.[1] Two cases motivate this position (Chalmers 2006, 53,107–8; Siegel 2010a, 169–70; Schellenberg 2011, 739–40). First, imagine seeing one object (say, a #2 pencil) and then swapping it with another that's not perceptibly different without changing anything else about the scene (lighting, your perspective, the other things around and their arrangement). Second, compare having a perceptual experience of an object (say that #2 pencil again) with an indiscriminable hallucination of it. In the first case it's intuitive to think that although the objects perceived are different the experiences are the same, while in the second it seems that you could not only have an hallucinatory experience that's indiscriminable, but one that is identical to the perceptual experience. Here I present a new, original argument for rejecting these intuitions and accepting that experiences are object dependent.

Conscious perception of an object allows you, by attending to it, to select it for thought. The argument I present starts with two claims about this perception-based thought. First, it's demonstrative. Second, experience of an object plays a role in it. That is, only *conscious* perception enables the selection of what's perceived for thought. I argue that if experiences are not dependent on the objects perceived in them, then at least one of these two claims is false. Although I won't develop new arguments for the two claims, both are highly intuitive, enjoy wide support, and have been argued for elsewhere. The upshot is a dilemma for those

---

[1]Throughout this paper I use 'perception' as a success term. Perceiving is always successful veridical perception, in contrast to hallucination or illusion.

who reject that experiences depend on the objects perceived in them: they must either reject that perception-based thought is demonstrative, or reject that experience plays a role in it. The dilemma is substantial. Perception-based thought is a pervasive and basic way in which you cognitively engage the world. Denying that it's demonstrative or that experience plays a role would require a radical rethinking of how it works.

The plan for this paper is as follows. The next section (§2.2) gives prerequisite background for perception-based thought. John McDowell and John Campbell have each also argued from perception-based thought to the object dependence of experience (Campbell 2002, ch 6,7, 2004; McDowell 1986, 156, 2010, 2013; see also Valberg 1992b, 21–2; Martin 2002a, 197–200; Brewer 2006, 179; Johnston 2006). Their basic line of argument is summarized in §2.3, along with a crucial reply from Tyler Burge. I conclude §2.3 by suggesting that Burge's reply fails to explain how experience plays a role. This point is a crucial new insight of the paper—the first—and motivates the idea that denying object dependence forces a dilemma. Next, §§2.4–2.6 turn my response to Burge into an argument for the dilemma, one that's independent of McDowell and Campbell's work. Two conditions are given (§2.4) that are jointly sufficient for the object dependence of experience: that all instances of an experience present the same object, and that an experience can be had only when the object presented in it is actually perceived. Then I argue that experience could not have a role in perception-based thought if the first condition failed (§2.5), while perception-based thought could not be demonstrative if the second condition failed (§2.6). Leveraging these jointly sufficient conditions is the second major insight of the paper.

It's important to note that the arguments from Campbell and others using perception-based thought purport to establish not just the object dependence of experience, but also naïve realism, the view that perceptual experiences are relational. But, as I explain below, whether experiences are object dependent is independent of whether they're relational or instead representational. Hence the argument given here is not for naïve realism and the conclusion that experiences are object dependent is consistent with representationalism (albeit perhaps not amenable to its motivations). The final section (§2.7) discusses these issues and explains how the argument falls short of establishing naïve realism.

2.2   Perception-Based Thought

Conscious perception allows you to select what's perceived for thought by attending to it (Pylyshyn 2001, 154, 2007, 18; Martin 2002a, 178; Johnston 2006, 265, 2011, 173; Siegel 2006a, 432, 2010a, 157; Speaks 2009, 560; Dickie 2010, 213–15; Jeshion 2010c, 132; Recanati 2010, 147,152; Brewer 2011, xii, 39–41; Smithies 2011b, 5; Wu 2011b, 109–11; cf. Hawthorne and Scala 2000). Consider the following example (compare with Smithies 2011b, 7).

> **Example 13:**
> On a table in front of you sits a standard tennis ball. You look directly at it (you foveate the ball and hold your attention on it) and judge that it's green.

In this example there's an occurrent mental act—a judgment—with some basic subject-predicate structured proposition as its content. You select the tennis ball as the *target* of the thought, i.e. as the subject to which the property *green* is predicated, by attending to it. More generally, call thoughts like this one in which you, by attending, exploit your conscious perception of an object to select it as the target *perception-based thoughts*.[2] Both occurrent propositional mental acts, such as judgments, and particular standing propositional attitudes, such as beliefs, can be perception-based. Note that this explication of perception-based thought by example assumes no account of propositional actions or attitudes (e.g., especially not a Fodorian LOT view).[3]

Example 13 is intended merely to point to cases in which attention to consciously perceived objects is used to select them for thought. Acknowledging perception-based thought doesn't require assuming anything about the nature of this selection. But it's very plausible and widely held that perception-based thought has the following two features. First, it's demonstrative, in some sense (Evans 1982; Campbell 2002; Martin 2002a, 179; Johnston 2006; Siegel 2006a, 432; Pylyshyn 2007; Speaks 2009, 560; Dickie 2010, 213; Jeshion 2010c, 133–35;

---

[2]Sometimes perception-based thoughts are called demonstrative thoughts (e.g. Smithies 2011b, 7), where 'demonstrative' doesn't refer to any special feature but merely denotes the ostensively defined category (that I've called) perception-based thought. I use the term differently; see immediately below.

[3]The example and highlighted phenomena of selecting consciously perceived objects for thought via attention is something that even an eliminativist about folk psychology could accept, perhaps with some minor reformulation. So, this isn't an "in house" or parochial debate. Whether experiences are object dependent is a deep, fundamental issue and the phenomena I leverage here to get at it, perception-based thought, is of similar fundamental status.

Levine 2010; Smithies 2011b; Wu 2011b, 109–11; Recanati 2012). Second, the experience of the perceived objects itself plays a role in enabling their selection (Valberg 1992b, 21; Campbell 2002, 7; Johnston 2006, 263–65; Siegel 2006a; Dickie 2011, 294,298; Smithies 2011a, 264, 2011b, 7, 19; cf. Kelly 2004, 283–4; Wu 2011b, 115–18). I shall argue perception-based thought can have both features only if experiences are object dependent. In the rest of this section I overview these two features.

That perception-based though is demonstrative is suggested by the fact that it's natural to express perception-based thoughts using a demonstrative term like 'that'. In example 13 you would naturally express your judgment by saying "*that* [pointing to the ball] is green". But in what sense, precisely, is perception-based thought demonstrative? One suggestion involves indexical representational vehicles: just as demonstrative speech acts involve demonstrative terms such as 'that', perception-based thoughts involve tokening analogous "mental demonstratives" (see Levine 2010; Recanati 2012). But this would be a very substantial claim (see Millikan 2012) and is not the premise I'll leverage in the argument for the object dependence of experience.

Before suggesting my preferred alternative, it's useful to say more about mental demonstratives. The view assumes a broader picture of thought on which selecting a target involves tokening a mental representation of that target. Mental representations, or their instances or realizers, are literal physical symbols in the head (e.g., Levine 2010, 181). A representation—not just a mental one—is *indexical* iff, for each token of that representation, its referent is determined by the context of tokening. For example, a token of 'that' refers to whatever is being demonstrated (attended, pointed at, etc) by the speaker. On the mental demonstrative view of perception-based thought, you select what's consciously perceived by tokening an indexical mental representation—called a *mental demonstrative*—the tokens of which have the functional role of referring to whatever is being attended.

The less controversial sense in which perception-based thought is demonstrative involves the capacity deployed. As Mike Martin points out (2002a, 178–81), selecting a target for thought by exploiting conscious perception involves exercising a one-off, episodic capacity.

> **Example 14:**
> There are many every day objects, properties, and locations with with you're acquainted, e.g. the car you drive, the town in which you live, and your friends. You think about them all the time without exploiting perception. You often think about them when you're

> not currently perceiving them. Right now, for example, I'm thinking
> about how my notebook is almost full.

When I have this thought I exercise a standing capacity to select my notebook. In contrast, when you select the tennis ball in example 13 you don't use any standing capacity. Instead, you select it in a way that depends on your current conscious perception of it. Standing capacities to select targets are underlied by concepts (see Peacocke 2009; see also Peacocke 1992b; Fodor 1998; Prinz 2002; Machery 2009). From my acquaintance with my notebook I have a concept, MY NOTEBOOK, and I select the notebook for thought by deploying that concept. This concept picks out a specific thing—my notebook—and (like all conceptual thought) the target of the thought deploying the concept is the thing picked out. Selecting a target by exploiting conscious perception is demonstrative in the sense that it doesn't involve deploying a concept.[4] Instead, conscious perception itself suffices for selection: you select what's perceived as a target for thought by *directly* exploiting your conscious perception of it.

Moving to the second feature, a view which enjoys some support is that experience plays a functional role in perception-based thought. That is, experience of the perceived object—conscious perception of it—is required, or necessary, to select it through attention. Alternatively (but equivalently) put, a perceived object can be selected for thought by attending to it only if that object is actually *presented* in experience. On this view, while perception is necessary it's not sufficient. There are two ways you might have perception without experience of the object. First, if (say) sense-data theory turned out to be correct, then although you perceive distal objects in the environment those objects wouldn't be presented in experience. What's presented—what you'd have experience of—would be sense-data. You would perceive the distal objects only indirectly. Second, assuming some form of direct realism is true, cases like blindsight in which there's some form of perceptual representation and processing, but no experience at all, would be cases of perception without experience.

To avoid misunderstanding, when I talk about what's *presented* in experience, I just mean the things of which you are aware in experience (Valberg 1992b; Fish 2009, 11), or the "objects of experience". Debates between direct realists and

---

[4]Episodic one-off capacities might be explained in terms of demonstrative or indexical concepts—mental demonstratives could be understood in this way—but this complication doesn't affect the point that perception-based thought doesn't involve mediation by redeployable concepts which refer to a unique object. Also, 'demonstrative concept' usually refers to concepts which underlie standing capacities, but have their referent fixed demonstratively (see Levine 2008).

sense-data theorists are, e.g., debates over what's presented in experience. Here I take direct realism for granted and assume that (at least in veridical, nonhallucinatory cases) the things presented are just the distal, mind-independent objects, properties, and relations that are successfully perceived in some sensory modality. Some philosophers use the term 'presentation' in a more loaded sense, e.g. making (by definition) something like Russellian acquaintance necessary for presentation (e.g., Brewer 2006, 169, 2011, 2; Martin 2002a, 173, 186; Johnston 2006, 267, 280). But I don't mean to burden the term with any substantial, stipulated conditions.

The main motivation for thinking that experience plays some role is that in both cases—indirect realism and blindsight—it's intuitive to deny the possibility of perception-based thought. First consider an indirect realist view like the sense-data theory. The usual intuition is that if you are presented with sense-data in experience, then it's only the sense data (not the objects indirectly perceived by having experiences of them) which are available for perception-based thought. Intuitive support from the direct realist perspective comes from considerations about super blindsighters and other cases of unconscious perception with access consciousness (e.g., Campbell's "sea of faces" case, 2002, 8). A super blindsighter cannot select an object that falls within the field defect (the "blind spot"), so the intuition goes, and the problem cannot be that she lacks access to visual content from that area or even that she doesn't perceive the object (Campbell 2002, 7; Johnston 2006, 264–265; Smithies 2011b, 6,26; cf. Kelly 2004, 283). The only difference between objects not in the blind spot (which the blindsighter can select) and those in it is that the blindsighter lacks awareness, or experience, of the latter. Hence, the inference goes, experience plays a role.[5] Aside from intuitions about blindsight, Campbell (2002) and Declan Smithies (2011b) have argued that experiences play a role by appealing to functions necessary for perception-based thought that, they argue, only it can fill.

2.3   Previous Arguments for Object Dependence

McDowell (1986) and Campbell (2002, 125; 2010, 201) both advance the following claim (see also Brewer 2006, 179; Johnston 2006, 264; compare with Adam Pautz's similar "grounding principle" 2009, 500).

---

[5]The claim that even super blindsighters cannot have attentionally select perceived objects is crucial, since without it one could raise Ned Block's classic concerns (Block 1995).

Basic Object-Dependency Thesis
Conscious perception allows you to select what's perceived for thought by attending to it only if experiences are object dependent.

They argue that were experiences not object dependent, then conscious perception would lack something required for it to enable selection. It would lack "directedness towards external objects" (McDowell 1986, 165) or would fail to "present" determinately a particular object (Campbell 2002, 125). As Martin noncommittally explains while elucidating McDowell and Campbell's insight (2002a, 198), perceiving one object obviously doesn't enable selection of other numerically distinct but perceptually indiscriminable objects. But if your experience when perceiving an object is the *same* as your experience when perceiving other objects which are perceptually indiscriminable, then—surely—conscious perception doesn't allow selection of the first even when that's the one actually perceived.

The above merely gestures at an intuitive puzzle for explaining how conscious perception allows selection of what's perceived for thought if experiences are not object dependent. Although they give some additional details (which I won't try to reconstruct), McDowell and Campbell do not go beyond this intuitive puzzle (see Burge 2005, 60,62; Pautz 2010, 283,286). The problem with leaving it there is that there's a plausible account of how conscious perception introduces what's perceived for thought—the mental demonstrative account summarized in §2.2—which, prima facie, doesn't require the object dependence of experience. So maybe the intuitive puzzle solves itself: once you get clear on how conscious perception enables selection you'll understand why experiences need not be object dependent. This response comes from Burge (2005, 44,54–63; see also Martin 2002a, 178–81,197).[6]

On Burge's response, perception-based thought involves the tokening of a mental demonstrative. Selection by mental demonstratives is a relational means of selecting. In relational means of selecting, whether conceptual or demonstrative, the target of thought is fixed by some relation between you (or concept, or mental demonstrative) and it. Mental demonstratives refer to whatever is attended, and the target of a perception-based thought is the referent of the mental demonstrative tokened in it. So on the mental demonstrative account, attention

---

[6]Martin (2002a) advances the discussion between Burge, McDowell, and (by extension) Campbell by setting out the issues more clearly than they had been before. My overview here of the basic object-dependency thesis owes much to this paper from Martin.

to consciously perceived objects enables their selection by relating you to them in a way you can directly exploit using mental demonstratives.

Now the crux: why think the attention relation is exploitable only if experiences are object dependent? Exploiting the attention relation just involves tokening a mental demonstrative with the functional role of referring to whatever is attended. There's no obvious reason why this requires that experiences depend on what's perceived. Perhaps the exploitability of the attention relation only requires that the consciously perceived (and attended) object is a constitutive part of the content of that particular instance of perceiving (Speaks 2009, 560; Siegel 2010a, 157; Schellenberg 2011, 739). Or, perhaps the involvement of causal relations in conscious perception is sufficient for the exploitability of attention to the consciously perceived things (Pautz 2009, 499, 2010, 286).

The way to respond to Burge, and these suggestions, is to press that even if perception-based thought is demonstrative in the strong sense of involving mental demonstratives (and hence is relational), if exploiting the attention relation doesn't require that experiences are object dependent, then it's unclear whether experience plays a role at all. The worry is that if selecting a consciously perceived object involves exploiting the attention relation to it (whether by tokening a mental demonstrative or not), but exploiting the attention relation doesn't require that the experience is dependent on the consciously perceived object, then the experience doesn't play a role in enabling selection. If the worry is well founded, then a relational account of perception-based thought can be used to save the object independence of experience from the pressure put on it by McDowell and Campbell only if experience doesn't play a role. Of course, all I've done so far is raise a worry in response to Burge. Nothing like an argument has been given yet (although see the end of §2.4, which develops the worry a bit further).

2.4   Two Conditions for Object Dependence

The last section concluded with the suggestion that if perception-based thought is relational but experience is not object dependent, then experience does not, or could not, play a role in perception-based thought. Since, as I argue below, the relational character of perception-based thought follows from it's demonstrative character (even construed more minimally as discussed above), this suggestion can be reformulated into the main thesis of this paper:

Main Thesis
If (a) perception-based thought is demonstrative, and (b) experience

plays some role in it, then experiences are object dependent.

As noted in the introduction (§2.1), this creates a dilemma for those who reject the object dependence of experience. They either need to give up that perception-based thought is demonstrative, or that experience plays some role in it.

In this section and the next two (§2.5 and §2.6) I'll argue for this thesis. The first step (this section) is to give two conditions which are jointly sufficient for the object dependence of experience. Before giving the two conditions, object dependence needs to be explained. An experience is *object dependent*, on a particular object O, iff, necessarily, that experience can be instanced only when actually perceiving O. An alternative way to understand object dependence is commonly found in literature on naïve realism and gets at object dependence by way of "constitutiveness". An object O is *constitutive* of an experience iff it's part of, or "shape[s] the contours of", the experience (Martin 2004, 64, cited in Fish 2009, 6; see also Campbell 2002, 116, Martin 2002a, 178). I take it that both approaches come to the same thing: what it is for an object to be constitutive of an experience is for the experience to be noncausally or necessarily dependent on it.

It's helpful to explain token experiences, since they provide a useful reframing of both object dependence and the two conditions. So far I've used 'experience' to refer to repeatable types, e.g. two occasions of perceiving might both instance—or perhaps involve (see fn 7)—the same experience (see, e.g., Siegel 2010a, 20). But you could also talk about particular instances of experiences. These *token experiences* are particular occurrent mental events (e.g., Fish 2009; Johnston 2011, 172–73; see also Pautz 2009, 494; Burge 2005, 34; Martin 2002a).[7] The *phenomenal character* of a token experience is the property that types it by "what it's like" for its subject to enjoy (Fish 2009, 8, who cites an unpublished work of Alex Byrne; see also Chalmers 2004, 341, 2006, 50). Experiences—or what I've been calling experiences so far—are just phenomenal character types of these token experiences. An experience (a phenomenal character type) is dependent on an object O iff, necessarily, the only token experiences which instance that type are ones in which O is perceived.

Turning now to the two jointly sufficient conditions, experiences are object dependent if:

---

[7]The metaphysics of token experiences won't matter here. For example, you might hold that token experiences are identical to the particular physical/causal events which are instances of perceiving, perhaps circumscribed to the mental aspects of those events (Byrne 2009, 431–35; Johnston 2011, 180), or instead you might hold that they're distinct sui generis events.

Condition 1: Presentation is phenomenally relevant: Necessarily, all instances of an experience present the same object.

Condition 2: Presentation is object dependent: Necessarily, an experience can be instanced only when the object presented in it is actually perceived.[8]

Both conditions assume the experience presents an object at all. Other kinds of experiences, e.g. an experience of a uniform color expanse, aren't at issue. The second condition assumes the first, since it takes for granted that there's a unique object presented in all instances of an experience. The need for the above restriction and the dependence of the second on the first can be eliminated by reframing the conditions in terms of token experiences: two token experiences which present different objects are not instances of the same experience, and a token experience presents O only if O is actually perceived. That the conditions jointly suffice for object dependence is straightforward on either framing.

Neither of the two conditions by itself suffices for the object dependence of experience. It should be clear that the object dependence of presentation isn't sufficient, since without phenomenal relevance it's left open that two token experiences might present different objects without instancing different experiences. But the phenomenal relevance of presentation isn't sufficient either. Consider the normal case of having an experience of one #2 pencil while perceiving it. Now consider two other cases. First, you have an hallucinatory experience of that #2 pencil. Second, you have an experience of that #2 pencil while actually perceiving a second that's indiscriminable. The phenomenal relevance of presentation neither precludes these cases as possibilities (as object dependence does), nor precludes the experiences had in these cases from being the same experience had in the normal case. So phenomenal relevance leaves open that the same experience could be had despite differences in what's perceived.

With the jointly sufficient conditions in hand, the argument for the main thesis has two steps:

Step 1: Argue that if presentation isn't phenomenally relevant, then experience doesn't play a role in perception-based thought.

Step 2: Argue that if presentation isn't object dependent, then perception-based thought isn't demonstrative.

If experience isn't object dependent, then presentation either isn't phenomenally relevant or isn't object dependent. So it follows that either experience doesn't

---

[8]Applying the terms 'phenomenally relevant' and 'object dependent' to presentation is merely a convenient way to refer back to the two conditions.

play a role in perception-based thought, or it isn't demonstrative. And, that's the main thesis. I give the needed argument for step 1 in the next section (§2.5), and the argument needed for step 2 after that (§2.6).

Before moving on, it's helpful to bring out the connection between the response to Burge from §2.3 and the two jointly sufficient conditions just given. The worry was that if perception-based thought is relational—involves selecting a target by exploiting some relation to it—but experience is not object dependent, then experience could not play a role in that selection. Pressing harder, the idea is that if perception-based thought is relational, then there must be some aspect of conscious perception that's object dependent, or relational. If not, then how could it enable a relational means of selection? But, if that very same aspect of conscious perception wasn't relevant to the experience—if experience was independent of the exploited object dependent aspect—then it seems as if there's no room for experience to play a role. The jointly sufficient conditions I've given here identify the experiential presentation of perceived objects as the aspect of conscious perception that's both object dependent and phenomenally relevant. Even if the arguments in the next two sections for the object dependence and phenomenal relevance of presentation fail, this point about there needing to be a single aspect of conscious perception that's both object dependent and phenomenally relevant would still stand. And, that point entails the object dependence of experience.

2.5   Step 1: From Experience's Role to Phenomenal Relevance

The argument required for step 1 is straightforward. Assume that presentation isn't phenomenally relevant: different instances of the same experience can involve perceiving different objects. Then if you select the object perceived in one of those two instances by attending to it, you would be selecting it in a way that's completely detached from the experience. That is, your selection of the object doesn't depend on using anything on which the experience depends. And if you can select the object in a way that is completely detached from experience, then it is just unclear how experience could play a role at all. Why wouldn't perceiving it indirectly (say, by experiencing some related sense-data) or perceiving it unconsciously (as in blindsight) suffice? The experience is in no way connected to the means by which you select the perceived object.

The main way to resist the argument is to reject an implicit assumption. The argument assumes that you select a perceived object for thought by attending directly to it. Before proceeding with alternatives, it's helpful to say more about

attention. The kind of attention through which a person selects perceived objects—I'm assuming—is personal-level voluntary (i.e., endogenous) focal attention. This is the kind of attention in play in familiar examples of attending, such as visually scanning a scene (which involves voluntarily, sequentially shifting attention between targets) or focusing on one sound (e.g. the speech of a particular person) while ignoring background noise (e.g. the racket of nearby traffic). It contrasts with involuntary (i.e., exogenous) attention, which is involved in cases in which your attention is captured by something salient in the scene (e.g., a sudden flash or loud bang), and with subpersonal attentional processes that underlie or bring about personal-level acts of attention (Allport 2011, 24–26). It also contrasts with nonfocal attention, i.e. attention that's divided between multiple things or spread around a larger chunk of the perceptual field.

Note that nothing hangs on assuming that the attention doing the selecting is voluntary or personal-level. The argument still goes through if it's subpersonal attentional processes which are involved (e.g. Campbell 2002; Clark 2006; Pylyshyn 2007; Levine 2010), or if involuntary attention works too. What matters is that the attending—however understood—is directly to the perceived object. Only if you select objects by attending to something other than the object itself is there a problem for the argument. To press the objection you might suggest that you select a perceived object by attending to one of its properties, parts, or features (or perhaps its location). One kind of account which might deliver this result is a sortal-based account on which you select a perceived object by attending to one of its properties, parts, or features while deploying a sortal concept. The selected object is whatever has the attended property, part, or feature and satisfies the sortal concept. In a similar way, any account which delivers this alternative would need to first invoke attention to a property, part, or feature F, then some further mechanism to bridge the gap between F and the perceived object.

But an account like this is implausible. Intuitively it does not seem to be the case that sortal concepts or any other further mechanisms are required for perception-based thought. If I have, say, a ball in clear view, then selecting it seems as simple as directing my attention towards it. For example, you can select the tennis ball in example 13 without having to think of it *as* a ball. And even if, as some have argued, selection requires the use of sortal concepts they typically aren't taken to play a role in bridging a gap between attention to a property, part, or feature of an object and selection of the object itself. For example, E.J. Lowe (2007) says that sortals are needed because selecting an object for thought

requires satisfying an epistemic constraint (requires "knowing which"), not because there is a referential gap to be bridged.

An even more basic problem for this strategy can be brought out by asking: Given that we can attend to objects, why would perception-based thought work in such a circuitous way, requiring that we direct attention away from the target to one of its properties? You might respond that, in fact, we cannot attend to objects themselves and can only attend to their properties, parts, and features. But this would be a very strong claim about attention, one that doesn't seem to be supported either by introspection or empirical research in psychology (Scholl 2001; Dickie 2010). Brain Scholl and his collaborators go so far as to say, in a paper working out the details of object-directed attention, that "that visual attention can select visual objects has thus been well confirmed, and has engendered many new theories of visual attention" (Scholl et al. 2001, 160).

The empirical research typically cited in support of object-directed attention includes the work on change blindness, on same-object advantages and "attentional spreading", and on multiple-object tracking (MOT) (see Scholl 2001, 5–13; see also Scholl et al. 2001; Scholl 2009). Most of the controversy in psychology over these results isn't about whether there is object-based attention, but instead about (a) the interactions between object-, feature-, and spatially-based attention and their relative priority in visual processing (Scholl 2001, 13–24), and (b) just what sorts of distal physical object-like things the visual system will treat as objects. Pylyshyn (who did much of the original MOT research) has expressed doubt over whether the "objects" this research shows are tracked are the distal, physical objects in the environment or instead are "proximal visual patterns or some other spatially local properties" of, or associated with, those objects (Pylyshyn 2003, §4.4.1,fn28, citation from anonymous reviewer). If his alternative is correct, then there is no object-based attention and what psychologists call object-based attention is just attention to "proximal visual patterns". But it's unclear what reason there is to think of attention as tracking proximal visual patterns (on the retina) of objects instead of the distal objects themselves. At the very least, the burden of proof seems to be on the person who resists the arguments being presented here. Given the empirical research is compatible with both views, why should we follow Pylyshyn's suggestion and think of it as showing that our attentional systems track proximal visual patterns and not the distal objects themselves?

2.6　Step 2: From Demonstrative Thought to Object Dependence

Recall that step 2 is to argue that if presentation isn't object dependent, then perception-based thought isn't demonstrative. The argument required for step 2 goes through three substeps.

Substep 1:　Argue that if perception-based thought is demonstrative, then it is relational.

Substep 2:　Argue that if perception-based thought is relational, then attention is also relational.

Substep 3:　Finally, argue that if attention is relational, then presentation is object dependent.

Substep 1:

Before giving the argument, it's worth saying more about relational means of selection. As noted already in §2.3, an episode (or state) of thinking, or the means of selecting involved in it, is *relational* iff the target of thought is fixed by some relation between you and it. Whether a thought or means of selecting is relational is a metasemantic question. Relational thought should be contrasted with descriptive thought. An episode (or state) of thinking, or the means of selecting involved in it, is *descriptive* iff the target of thought is fixed via satisfaction. In descriptive thought some condition is set out—perhaps a condition encoded in a descriptive concept employed in the thought—and the target is the unique thing which satisfies that condition. What's important is that having a relational thought requires that you stand in the relevant relation to the target. In contrast, aside from satisfaction no particular relation needs to hold between you and the target in descriptive thought; the target need not have ever existed.

　　If perception-based thought is demonstrative in the sense that it involves mental demonstratives, then step 1 is trivial. On the mental demonstrative account of perception-based thought, the referent of a given tokening of the mental demonstrative is determined by context. It's referent is the thing to which you're attending at the time of tokening—or whatever the contextual relation happens to be. Since, on this account, the target of a perception-based thought is the referent of the token mental demonstrative involved in the episode of thinking, it follows immediately that perception-based thought is relational. Through the use of mental demonstratives, the target of a perception-based thought is fixed by the contextual relation to which the mental demonstrative is sensitive. Any

mental demonstrative account will identify some relation which fixes the referent of token mental demonstratives. So, any mental demonstrative account will make perception-based thought relational.

But, as noted above, the mental demonstrative account is a substantial position. The main thesis takes perception-based thought to be demonstrative in the second, less controversial sense given in §2.2. Even this way of being demonstrative—involving one-off, episodic capacities to select—supports the claim that perception-based thought is relational. While on the mental demonstrative account being demonstrative entailed, by definition, being relational, the inference from perception-based thought involving one-off, episodic capacities to it being relational is an inference to the best explanation. Being relational itself wouldn't explain why conscious perception affords a one-off capacity to select. For example, the thought in example 14 is relational, but involves a standing capacity to select. But, a means of selecting which involved exploiting a relation directly—not by deploying a concept the referent of which is fixed by a relation—would explain it. If you exploit the relation directly, without deploying a concept the referent of which is fixed by that relation, then you have the capacity to select only as long as the relation holds. There is a way to explain the one-off capacity by appealing to descriptive concepts. For example, using the concept OBJECT CURRENTLY PERCEIVED would give you a capacity to select what you're currently perceiving only as long as it's being perceived. But this explanation isn't as good as the relational explanation because it violates the intuitions (mentioned in §2.5) that perception-based thought doesn't involve the use of concepts.

Substep 2:

The last substep argued that if perception-based thought is demonstrative, then it's relational. This substep argues that if it's relational, then attention is too. Recall from §2.5 that it's personal-level voluntary focal attention at issue. The claim that it's relational is the claim that the acts or episodes of voluntary focal attention in which subjects engage are object-involving. You cannot, for example, attend to items in the scene around you as you do visual search without those items actually being there and you causally engaging with them in the right way. Voluntary focal attention essentially involves the attended objects. As with the caveat in §2.5, although I assume it's personal-level voluntary focal attention involved in perception-based thought, the arguments in this substep go through even if subpersonal attentional processes or involuntary attention are what's involved.

Assuming perception-based thought is relational and that in it you select the

target by attending to it, a natural move is to infer that attention provides the exploited relation which fixes the target. (This was assumed in §2.3, for the sake of exposition.) If perception-based thought is relational, then selection happens by exploiting some relation. It's difficult to see how attention could not be the exploited relation, but still play an integral part in selecting the target. Alternatively put, if attention is not relational and you select the target in a perception-based thought by attending to it, then perception-based thought isn't relational.

There are accounts of perception-based thought on which it's relational, but attention isn't the exploited relation. For example, some versions of the mental demonstrative view take mental demonstratives to be, or be labels for, mental files (Bach 1987; Raftopoulos and Müller 2006, 256–58,264; Dickie 2010, 224, 2011, 304; Jeshion 2010c, 129–35; Recanati 2012, 34–38). A *mental file* is functionally defined as a collection of information or beliefs which our cognitive processes treat as being about the same thing. These mental files have as their referents the dominant casual source of information in the file. One possible view is that there's a special-purpose temporary file that's opened to store incoming perceptual information and closed (and reopened) as what's being perceived changes. This file is the mental demonstrative involved in perception-based thought. The point to note is that the mental file account isn't a problem. For the mental files tokened in perception-based thought, the dominant causal source will be the object attended. Attention, plausibly, will be a subrelation constitutive of the causal relation which fixes the referent of the mental file. If not, then (again) it's unclear how or why attention is an integral part of selection. The result is that attention, thanks to its role in selecting the target, must be relational.

Substep 3:

The final substep in step 2 is to argue that if attention is relational, then presentation is object dependent. This move turns on a widely held view on the relationship between attention and conscious perception. On this view, often only tacitly assumed, attention operates over what's consciously perceived, serving as a mechanism for selecting from among those things (Valberg 1992b, 21; Scholl 2001, 20; Pylyshyn 2007, 59; Levine 2010, 181; Dickie 2010, 216, 2011, 303 Wu 2011a, 109). Importantly, on this view attention and conscious perception are distinct mental acts or episodes, perhaps underlied by functionally independent subpersonal processes. Objects are consciously perceived (presented in experience), and then attention is a distinct process which operates over those objects. The alternative view takes conscious perception and attention to be more inte-

grated, both at the level of mental acts or episodes and the underlying supersonal processes. Accounts on which attention is necessary for conscious attention are one example (e.g. Prinz 2011).[9]

On this view of the relationship between attention and conscious perception, if presentation is not object dependent, then attention is not relational. Imagine that you are having just the experience you did in the tennis ball example (example 13), but this time are hallucinating that tennis ball. The case is possible if presentation is not object dependent. Now, assuming that attention operates over what's consciously perceived, there's no principled reason to prevent it from operating in just the same way over all experiences, even hallucinatory ones. So, attention to the tennis ball should be possible when it's only hallucinated in precisely the same way as when you perceived it. But in that case attention is not relational: attention to an object wouldn't require being properly related to it. In general, if conscious perception—the presentation of objects in experience—has the role of grounding attention by providing the items it selects, then object independence in presentation will break the connection between attended objects and acts or episodes of attending which made attention relational.

You might object that attention could be relational while presentation is object independent, even assuming that attention operates over what's consciously perceived, if conscious perception had relational aspects which did not affect experience, but were still required for attention. Loosely, perhaps attention requires both nonrelational mental components such as presentation and relational, nonmental components such as causation. This suggestion is in line with how Pautz responds to similar arguments from demonstrative thought for disjunctivism (2010, 286). It's a version of a general strategy, already mentioned (§2.3), of explaining how conscious perception could enable a relational form of thought without experience being object dependent. It's unclear how to respond to this objection only drawing on (a), the assumption that perception-based thought is demonstrative. But falling back to the original response with which §2.3 concluded, if the relational aspect of conscious perception exploited in perception-based thought is detached from experience in the way this objection proposes, then how experience could play a role is left unexplained. The view of perception-based thought given in the objection would imply that nonconsciously perceived objects can be selected for thought. So blocking this objection to this substep of

---

[9]Some versions of the alternative would still support substep 3. Jesse Prinz (2011), e.g., would deny that attention selects what's presented in experience. But, he *identifies* being presented with being attended (2011, 175), making step 3 trivial.

the argument (substep 3) requires appealing to (b), experience having a role in perception-based thought. But this doesn't weaken the overall argument for the main thesis, of course, since (b) is already one of its antecedent conditions.

2.7 Naïve Realism

The claim that experiences are object dependent is a central part of naïve realism. For example, Heather Logue (2012b, 221) describes naïve realism as "the view that veridical experience fundamentally consists in the subject perceiving things in her environment and some of their properties". Gordon Knight says (2013, 1), "On naïve realism, the phenomenal character of visual experience is accounted for by the qualities exhibited by the physical objects of which one is aware." Kenneth Hobson (2013, 551) similarly says that on naïve realism "veridical experiential states are dependent metaphysically on their objects for the reason that such objects are constituents of experiential states" (see also Campbell 2002, 116; Martin 2004, 39; Brewer 2011, 92; Johnston 2011, 181; Genone forthcoming).

But naïve realism goes beyond the claim that experiences are object dependent. It says that they are relational states; having an experience of an object is a relation. It contrasts with representationalism, the view that experiences are representational states (see Crane 2006).[10] Importantly, experiences could be representations and still be object dependent, showing that the argument in §§2.4–2.6 doesn't rule out representationalism (the point is well known, but see Genone forthcoming). More importantly, there's no straightforward way to extend the argument to rule out representationalism. More details are needed to distinguish fully representationalism and naïve realism (see Genone forthcoming), but they aren't needed to explain why the argument can't be extended.

Before discussing why the argument doesn't extend, consider how experience could be object dependent on representationalism. A version of representationalism will make experiences object dependent if, on it, experiences represent the particular objects presented in them and do so in an object-dependent way, i.e. if they have object-dependent content or are object-dependent repre-

---

[10]Since by 'experience' I mean phenomenal character types, this is intentionalism, the view that there's no change in phenomenal character without a change in representational content. This is an accident of exposition. Reframing in terms of token experiences would avoid intentionalist commitments: token experiences are representational mental states or episodes. The representationalists I cite below aren't necessarily intentionalists. The exposition here of representationalism doesn't bring out its nuances. They could be brought out by adopting talk of token experiences, phenomenal character, content, gappy content, (and so on). But doing so wouldn't contribute to the few and relatively simple points of this section.

sentations (Speaks 2009, 554 describes this view). Fully working out this version would require saying how the content was object dependent. Is it Russellian or Fregean with de re modes of presentation? It would also require a metasemantic account of how experiences got their content. Presumably on this account experiences are relational representations, i.e. have their content fixed relationally. Their relational character would explain the object dependence of their content. What matters is that any version of representationalism on which experiences are object dependent representations of the particular objects presented in them will make experiences object dependent.

An object-dependent version would go against one of the main motivations for representationalism, which is to the save the intuition that in-principle indiscriminable experiences are the same (e.g. Schellenberg 2011, 739–40). For example, representationalism provides a way to be a direct realist (as opposed to, e.g., a sense-data theorist) while respecting the intuition that the experience you enjoy while perceiving an object is the same as in an indiscriminable hallucination of it.[11] Representationalists avoid object dependence in three ways. First, they can deny that experiences represent particular objects and attribute only general or existential content to them (e.g. Davies 1991, 1992; Tye 1995; Pautz 2009; see also Searle 1983). Second, they can make experiences represent particular objects, but make them (or the aspect which represents the particular object) indexical or demonstrative-like representations. Susanna Schellenberg's gappy-content view is an example of this (2011). It's not that each experience represents a particular object. Instead, experiences are like indexical terms, a given instance of an experience representing what's perceived. Third, they can make experiences represent particular objects nonindexically, but still take for granted that tokening the experience didn't depend on actually perceiving the object it represented. (Other examples of the second and third positions include Soteriou 2000; Dretske 2003; Chalmers 2004, 2006, Speaks 2009; and Burge 2010.) Although most representationalists adopt one of these three, the basic commitment that experiences are representations doesn't rule out object-dependent versions.

Why won't the argument given here extend to rule out object-dependent versions of representationalism? First, note that these versions of representational-

---

[11]Representationalism also provides an explanation of illusions that many find appealing, viz. that they are experiences which misrepresent. But a version of the view which makes experiences object dependent could give this same explanation: making representations of the perceived object dependent on it doesn't imply that the properties attributed to it in those representations likewise depend on the actual properties.

ism satisfy both conditions 1 and 2, they make presentation phenomenally relevant and object dependent. To get naïve realism, the second condition needs to be strengthened (see Fish 2009, 14–16).

Condition 2∗**:** Presentation is relational: Having an experience present an object doesn't involve representing that object; instead, it involves standing in a sui generis relation to the object.[12]

Step 2 of the argument argued that the demonstrative character of perception-based thought requires that presentation is object dependent. The reason the argument can't be extended is because the demonstrative character of perception-based thought does not require that presentation is relational.

To see why presentation need not be relational for perception-based thought to be demonstrative, recall that the argument in step 2 had three substeps. Substep 3 was to show that if attention is relational, then presentation is object dependent. This substep can't be extended; it's not the case that if attention is relational, then presentation is relational. To make this clear, imagine that an object-dependent version of representationalism is the correct view of experience. Then, although the presentation of an object in experience is a matter of representing the object and not a relation, the way of representing itself essentially involves the object. Assuming that attention operates over, or is a means of selecting, what's presented in experience, this suffices to make attending to an object an act that involves the object, i.e. to make it relational. There couldn't be cases, like the one given in the argument in substep 3, of attention to an object when it's not actually being perceived.

Stepping back from the details of the arguments given here, experiences are object-dependent representational states on object-dependent versions of represenationalism. The crux is that object-dependent representational states are relational, at least in the sense that they are object-involving. And that is all that's required for perception-based thought to be relational, that experiences are relational in the sense of involving the objects perceived, i.e. being object dependent.

---

[12]Presentational could be relational without it being phenomenally relevant. As before, the conditions could be reframed in terms of token experiences to make this independence clear.

Chapter 3:
Perceptual Links in Demonstrative Thought

3.1   Introduction

Conscious perception allows you to select the things you perceive for thought (Pylyshyn 2001, 154, 2007, 18; Martin 2002a, 178; Johnston 2006, 265, 2011, 173; Raftopoulos and Müller 2006, 253; Siegel 2006a, 432, 2010a, 157; Speaks 2009, 560; Bach 2010, 55; Dickie 2010, 213–15; Jeshion 2010c, 132; ; Levine 2010; Recanati 2010, 147,152, 2012, 12,29 ; Brewer 2011, xii, 39–41; Smithies 2011b, 5; Wu 2011b, 109–11; cf. Hawthorne and Scala 2000).[1] More specifically, it *relates* you to what's perceived in a way that can be exploited to select the consciously perceived thing as the (variously called) target, subject, or referent of a thought. Call thoughts in which you exploit this perceptual relation, or *perceptual link* (Dickie 2010, 213, 2011, 297, forthcoming), *perception-based thoughts*. Perception-based thoughts are often called demonstrative thoughts, but I avoid this name because it suggests they share features with demonstrative speech acts when this does not need to be assumed (see §3.2.1).

This raises the question, what is the relation exploited in perception-based thought (Pylyshyn 2007, 97)? A natural suggestion is that the relation is conscious perception itself. Consciously perceiving a thing is a way to be related to it, and this relation is the exploited perceptual link. But this won't work, for the simple reason that generally you are consciously perceiving multiple things but are still able select one of them for thought (Levine 2010, 178). You could not exploit a relation holding between you and multiple things to select one of them; it would leave the selected one underdetermined. In this paper I argue for a new answer to this perceptual link problem: the exploited perceptual link is the rela-

---

[1]It's widely assumed that *conscious* perception is required for selection (Valberg 1992b, 21; Campbell 2002, 7; Johnston 2006, 263–65; Siegel 2006a; Dickie 2011, 294,298; Smithies 2011a, 264, 2011b, 7, 19; cf. Kelly 2004, 283–4; Wu 2011b, 115–18). Although nothing to follow depends on it, I frame my discussion in these common terms.

tion of voluntary focal attention, i.e. the relation in which you stand when you voluntarily attend to a single consciously perceived thing.

Specifically, in this paper I articulate this voluntary focal attention view (§3.3) and provide motivation for it (§3.3.2). I then argue it works better than its two competitors: a more permissive attention-based view (§3.4) and the received dominant causal source view (§3.5). The motivation is that intuitive assessments of the target in cases of perception-based thought follow voluntary focal attention. The permissive attention-based view lacks the restriction to *voluntary* attention and, similar to conscious perception itself, faces a fatal underdetermination problem. The received dominant causal source view presupposes that thinking involves mental files (a view widely held in philosophy and psychology) and takes the perceptual link to be the relation in which you stand to the dominant causal source of information in the relevant mental file. Here the problem is that there are cases of perception-based thought in which what you'd intuitively assess as the the target is not the dominant causal source of information (but is what's voluntarily attended). Note that it's not my purpose to give a definitive positive argument for the voluntary focal attention view or to defend it against all possible objections. Also note two other contributions of the paper. First, §3.2 opens the paper with a clear articulation of the perceptual link problem, something not found in the literature. Discussion of the problem either builds in substantial assumptions about the relational and demonstrative character of perception-based thought (§3.2.1) or doesn't distinguish it from broader questions about information channels (§3.2.2). Second, I clearly distinguish the voluntary focal attention and dominant causal source views, again something not found in the literature.

Before starting the initial overview, note one way the perceptual link problem is important.[2] Perception-based thought is often taken to play a foundational role in our cognitive lives. First, it's a way to have relational, or singular, thought (see §3.2.1). Second, it's a way of selecting targets for thought without having concepts of them (Campbell 2002; Raftopoulos and Müller 2006, 252–3; Pylyshyn 2007, 56). Third, it's needed for acquiring concepts through experience (Putnam 1975; Raftopoulos and Müller 2006, 261; cf. Levine 2010, 193). Learning the concept RED, for example, requires seeing instances of red, exploiting conscious perception to select the red, then categorizing that selected bit of red as an instance of red. Overall, these three points are often taken to make perception-

---

[2]There are other ways too. For example, a worked-out answer will surely inform work on naturalizing representation (e.g., Dretske 1981, 1988, 1995; Millikan 1984, 2004; Fodor 1987; see also Pylyshyn 2007, 97).

based thought, along with conscious perception itself, a primitive kind of intentionality which grounds derivative forms of mental and linguistic representation (e.g., Campbell 1997, 55; Pylyshyn 2007, 19; Levine 2010, 169). Crucially, many questions about precisely how and the extent to which perception-based thought can play these roles turns on the exploited perceptual link. For example, you might think that perception-based thought allows for having nonconceptual, relational thought about particular objects, or that it allows for learning certain high-level concepts (e.g., CHAIR) through experience. But whether it allows for either depends on whether perceptual links go to particular objects or high-level properties, and that depends on the identity of the exploited perceptual link.

3.2   The Perceptual Link Problem

Perception-based thought can be introduced by example (compare with Smithies 2011b, 7).

> ### Example 15:
> On a table in front of you sits a green tennis ball. You look directly at it (you foveate the ball and hold your attention on it) and judge that *that's* green.[3]

The thought in this example predicates a property to a subject. That is, it has basic subject-predicate structured content. Call the tennis ball, the subject to which the property *green* is predicated, the *target* of the thought. Intuitively, the way you select the tennis ball as the target has something to do with your conscious perception of it. Contrast the selection in example 15 with a second case in which perception isn't directly involved.

> ### Example 16:
> You move away so the ball is out of sight. Thinking about it (e.g., judging that it's green) now requires remembering it, deploying a concept of it you acquired while looking at it, or giving a description which picks it out.

A plausible suggestion is that in the perception-based case (example 15) conscious perception grounds a relation to the tennis ball, a perceptual link (as Imogen Dickie calls it, 2010, 213, 2011, 297, forthcoming), which you exploit to select

---

[3]Although the example involves selecting a particular object, examples in which properties of objects, or locations, are selected would equally work. I use the term 'thing' as a neutral way to refer to potential targets of selection.

it. What makes the ball the target is that it's the thing at which you're looking, or rather, it's the thing to which you have a perceptual link. This leads to the perceptual link problem: what is that relation you exploit to select the target?

An immediate point of clarification is that selecting is a metasemantic operation: thoughts have targets and something must fix, or select, a thing as the target (Pylyshyn 2007, 3–8,23; Levine 2010, 172). To elaborate, compare thoughts with assertions (the uttering of declarative sentences in speech acts with assertive force) which also have targets which must be selected. Typically the target of an assertion is selected by using a term in the spoken sentence which refers to that target (e.g., a name, description, anaphoric pronoun, or indexical). An account which similarly explains selection in the case of thought won't be required here.[4] (A typical one which explains selection in terms of concepts, or mental representations, is given in §3.2.1 as part of the exposition of the dominant causal source view; it's not presupposed by the perceptual link problem or the voluntary focal attention view.) But one potential confusion should be avoided. While conscious perception also (a) makes available concepts of what's perceived and (b) often provides justification for thoughts (e.g., beliefs and judgments) about what's perceived, neither should be confused with how it allows for selection. Making available a concept of a thing and providing justification for thoughts about it are not the same as providing a means to select it as a target.[5]

3.2.1   Relational and Demonstrative Thought

Because discussions of perception-based thought often assume that it's relational and demonstrative, it's important to point out that the perceptual link problem can largely be broken free of these assumptions. First, the exploiting of a perceptual link to select a target is consistent with perception-based thought being either relational or descriptive. A thought is *relational* if its target is fixed by a relation between the person who has the thought and it. It is *descriptive* if its target is fixed by satisfying a condition. The idea is that in some thought the target isn't picked out "directly" by a relation to it, but instead there's a descriptive condition (e.g., *the tennis ball on the table*) and the target is the unique thing satisfying that condition (Jeshion 2010a, 1; Recanati 2012, 3–6; Dickie forthcoming, §2). The re-

---

[4]Talk of "exploiting" a relation, as I use here (see also Recanati 2012, 21), is just a stand-in for this missing account.

[5]You can select an object as a target by deploying a perceptually acquired concept of it. (Deploying concepts in thought is like using words in speech acts, see §3.2.1.) Even if this happens in example 15, the selecting would be something distinct from the acquisition of the concept. But it doesn't seem to be what happens (see §3.2.1).

lational/descriptive distinction is best known from the debate over reference in philosophy of language between direct reference theorists (Barcan Marcus 1961, 310; Kripke 1981) and descriptivists (Searle 1983).[6] But a similar distinction holds for thought, where often relational thought is called singular thought. (See Jeshion 2010b for recent work.)

How is the perceptual link problem consistent with either view? It's natural, and standard (e.g., Campbell 2002; Raftopoulos and Müller 2006, 253,261–2; Bach 2010, 55; Jeshion 2010c, 134; Recanati 2012, 12,37; Dickie forthcoming), to understand exploiting a perceptual link in relational terms, so that perception-based thoughts are relational. Exploiting a perceptual link involves the link directly fixing the target of thought, without involving the satisfaction of a condition. For simplicity, I assume this relational account below. But a descriptivist explanation of perception-based thought will also appeal to perceptual links, albeit in a different way. For example, a natural descriptive account of perception-based thought is that the target is whatever satisfies some condition equivalent to *the thing to which there's currently a perceptual link*. The perceptual link comes into play as a component of the descriptive condition, and so the descriptivist still must answer the perceptual link problem.

Moving to demonstrative thought, while perception-based thought is often taken to be demonstrative (Evans 1982; Campbell 2002, 114; Martin 2002a, 179; Johnston 2006; Raftopoulos and Müller 2006; Siegel 2006a, 432; Pylyshyn 2007; Speaks 2009, 560; Dickie 2010, 213; Jeshion 2010c, 133–35; Levine 2010, 169; Smithies 2011b; Wu 2011b, 109–11; Recanati 2012) the perceptual link problem only presupposes this demonstrative character in one noncontroversial sense. The demonstrative character of perception-based thoughts is suggested by the fact that they're most naturally expressed with demonstrative terms like 'that' (e.g., Levine 2010, 170–1). For example, your judgment in example 15 would naturally be expressed by saying "*that* [pointing to the ball] is green". There are at least three things which might be meant by saying that perception-based thought is demonstrative. The first is that it's "minimally direct", in the sense that it's not descriptive thought involving properties of the consciously perceived thing. The second is that it involves exercising a one-off, episodic capacity to select the consciously perceived thing. The third is that it involves the token-

---

[6]An importance difference is that the debate over names is usually cast as all or nothing: either names directly refer, or have descriptive contents. For thoughts, everyone agrees that some thoughts are descriptive. The debate is over whether any are relational. Perception-based ones are the best bet, but the perceptual link problem doesn't presuppose they are relational.

ing of demonstrative-like mental representations, "mental demonstratives". The perceptual link problem presupposes only the first of these, but explaining all three will help elucidate perception-based thought and provide background for the dominant causal source view.

Although perception-based thought involves properties of what's perceived insofar as it involves predicating properties to it, the properties of the perceived thing do not get involved in the selection of it as the target (Campbell 2002, 7; Raftopoulos and Müller 2006, 252–3,264; Pylyshyn 2007, 14–6; Jeshion 2010a, 1, 2010c, 134; Levine 2010, 185; Recanati 2012, 29). Consider the following way you could use your conscious perception in example 15 to select the tennis ball, but is not how the selection happens when you exploit your perceptual link. You might first notice features of the ball, like it's shape, kind, and location, then descriptively select the ball as target to which *green* is predicated by thinking of it in terms of those noticed features. You might select it by thinking of it *as* the round tennis ball on the table, i.e. by using the descriptive condition *the round tennis ball on the table*. But your conscious perception of the ball provides a link which allows you to bypass this describing step and select the ball directly, just by exploiting your perceptual link to it. Call this *minimal* directness, since even if perception-based thought turned out to be descriptive, the descriptions would not involve nonrelational properties, instead involving conditions which refer to the perceptual link. The perceptual link problems assumes that perception-based thought is demonstrative in this sense.

Moving to the second sense, when you consciously perceive something, the means which it affords you to select that thing as a target depends on you actually perceiving it. Although the perceptual link problem doesn't presuppose it, it's plausible that perception-based thought is demonstrative in this sense. As Mike Martin notes (2002a, 178–81; see also Recanati 2012, 62), conscious perception gives you a one-off, episodic capacity to think about what's perceived. The relevant contrast is with cases that involve using a standing capacity, such as thinking about the tennis ball when it's no longer in view (example 16). Typically, a standing capacity to select a particular thing as a target of thought is underlied by memories of past experiences of it, singular or acquaintance-based concepts you have of it, or the use of descriptive concepts—concepts which "encode" descriptive conditions—that pick the thing out. Concepts are often said to be constituents of thought, sometimes understood as physically tokened, semantically typed representations in the head, sometimes as abstract constituents of the propositions which are the contents of thoughts (see Peacocke 2009; see also

Peacocke 1992b; Fodor 1998; Prinz 2002; Machery 2009). On the former view (Raftopoulos and Müller 2006; Jeshion 2010c, 129–35; Levine 2010, 169–70; Recanati 2010, 181, 2012, 35–7,64; Taylor 2010, 83–7; Dickie 2011, 292,297; see also Pylyshyn 2007),[7] deploying a concept is tokening a representation in the head, thinking consists of tokening representations in certain ways, and selecting a thing as a target involves tokening a representation (a concept) of that thing. On this view there's a rough analogy between how concepts function in thinking and how words are used in speech acts.[8]

The third sense in which perception-based thought is taken to be demonstrative is that it involves indexical mental representations called mental demonstratives (Bach 2010, 55; Levine 2010, 179; Recanati 2012, 57–67). This view presupposes that selecting a target involves tokening a representation of it, but further posits that selecting in perception-based thought involves tokening mental representations which, like demonstrative terms (e.g., 'that'), are indexical. A representation is *indexical* iff, for each token of that representation, its referent is determined by the context of tokening. For example, a token of 'that' refers to the thing being demonstrated (attended, pointed at, etc) by the speaker. Most versions of the dominant causal source view presuppose that perception-based thought involves mental demonstratives, since the mental files they presuppose either act as, or are labeled by, mental demonstratives.

3.2.2 The Information Channel Approach

Starting with Gareth Evans (1982), perception-based thought has largely been approached in terms of information channels (Recanati 2010, 156, 2012, 35; Taylor 2010; Dickie 2010, 2011, forthcoming). An *information channel* is any relation in which you stand to something which allows you to gain information about it, typically but not necessarily causal (Millikan 2004, 44). What motivates interest in information channels is that they, whether perception-based or not, are often thought to be exploitable as means of selection, allowing for relational thought. This fact, plus the fact that conscious perception is a paradigmatic information channel, is plausibly what's lead to framing perception-based thought in these terms.

---

[7]Often the term 'concept' isn't used and the view is framed simply in terms of mental representations or mental files.

[8]Despite the analogy, the view that thinking involves physically tokened, semantically typed representations in the head isn't automatically Jerry Fodor's language of thought (LOT) view (Fodor 1975, 2008). A LOT view would attribute specific properties to these representations, e.g. a syntax and semantic constitutionality.

There are two ways to understand the information channel framing. The first begs the question and assumes that the answer to the perceptual link problem is some relation defined in information-theoretic terms. For example, the relation which holds between you and a consciously perceived object in virtue of you having a perception-based information channel might be identified as the exploited link (e.g., Raftopoulos and Müller 2006, 264). The dominant causal source view is one variation on this view. Within a discussion of the perceptual link problem such an approach should be avoided.

The second way only assumes that, whatever the perceptual link, it must be an information channel. This leaves room for identifying the perceptual link in noninformation-theoretic terms, e.g. as the voluntary focal attention relation. This is the widely accepted acquaintance condition on relational thought (see Jeshion 2010a, 14–16, e.g. Bach 1987; Dickie 2011; Recanati 2010, 2012; Genone 2014). The condition, based on Russell's notion of acquaintance (1997/1912), says that relations exploitable for selecting targets directly need to acquaint the thinker with the target. I think this condition is reasonable (although see Jeshion 2010c). It still allows for a neutral framing of the perceptual link problem. Since voluntary focal attention is an information channel, the voluntary focal attention view is not ruled out by the condition. The point is just that we can accept the acquaintance condition without actually defining the perceptual link in information-theoretic terms.

3.3   The Voluntary Focal Attention View

My proposed answer to the perceptual link problem identifies the exploited perceptual link as voluntary focal attention. In example 15, to take one case, it's your voluntary attention to just the tennis ball which fixes it as the target of thought. Your current focused, voluntary attention to the tennis ball relates you to it, and it's in virtue of this relation that the ball is the target of thought. Attention-based views like this one originate with Campbell and are common (Campbell 1997, 2002, 2004; Levine 2010; Wu 2011a,b), but the distinction between voluntary and involuntary attention is typically not addressed. This section explains voluntary focal attention (§3.3.1), then gives preliminary motivation for identifying it as the exploited perceptual link (§3.3.2). Because Campbell's view is so prominent, §3.3.1 concludes by contrasting it with the voluntary focal attention view.

3.3.1   Voluntary Focal Attention

Attending is an action directed at consciously perceived things. It's what you do, e.g., when asked to attend to, or focus on, some perceived object, location, event, sound, smell, feel, pain, or other consciously perceived thing. Alan Allport (2011, 25) aptly describes it as "a state or relationship of the *whole organism* or person", and says that "Prototypically, we are talking about a transient state in which a person's coordinated, purposefulx thought and action can be directed toward, or guided by, the object of attention." Moving from rough-and-ready descriptions like this one to specific analyses is difficult. (See Carrasco 2011 and Mole et al. 2011a for recent work.) There are two approaches: focusing on the functional role of attention, and focusing on how attention affects experience, i.e. the phenomenology of conscious perception. For example, whether the view defended here on the role of attention in perception-based thought is correct, it's widely thought that attention functions in cognitive tasks as a means of selecting (e.g., Campbell 1997, 2002; Pylyshyn 2007, 14,57; Levine 2010, 178; Carrasco 2011, 1486; Wu 2011a). And, at the level of experience attention seems to divide the perceptual field into a foreground and background (Watzl 2011) and increase the apparent contrast of what's perceived (Carrasco et al. 2004; Wu 2011b).

For the purpose of explaining the voluntary focal attention view, it suffices to give examples of attending. Four examples include: focusing on one of two simultaneously playing audio streams (e.g., a bit of speech and background noise), visually tracking an object as it moves past you, visually searching a scene for a specified object, and being startled by a loud bang or bright flash. In the first case you hold attention to one audio stream, in the second case you hold attention on the moving object, in the third you shift attention around the scene, and in the last case attention is grabbed by the sudden noise or flash.

Attention can be focused on one thing, or divided between multiple things.[9] For example, you might try to attend simultaneously to both an audio stream and visually track a moving object. Or, you might try to visually track multiple objects at once, or track a moving object while loud bangs and bright flashes continuing grab your attention. For a final example, pick a word on this page and shift your attention between it and the word to its left, then attend to both words at the same time. Any attention-based answer to the perceptual link problem must be limited to focal attention for the obvious reason that divided attention would leave the target underdetermined.

---

[9]I have borrowed the term 'focal attention' from Pylyshyn (2007, 59).

Attending, as I'm using the term, should be distinguished from the functionally defined representation-manipulating processes carried out in the brain and the neurobiological mechanisms that enable, "[bring] about", or explain personal-level acts of attention (Mole et al. 2011b, xi; see also Allport 2011, 26; Watzl 2011, 147; Wu 2011a, 106, 2011b, 97). These "subpersonal" attentional processes are mechanisms that pick out information, or representations that carry information, transduced from the sensory organs for further processing. It's a substantial part of the voluntary attention view that the exploited link in perception-based thought is the "personal-level" act of attending and not one that holds in virtue of the operation of these subpersonal processes.[10] (For example, Levine 2010 suggests the latter view.) I won't defend it here. Instead I focus on the distinction between voluntary and involuntary attention. A defender of the subpersonal version of the view could suitably adapt the arguments to follow into a defense of her view over versions that allow involuntary attention and over the dominant causal source view.

Attending can be voluntary (endogenous) or involuntary (exogenous). The examples given above are the starting point: in the first three it seems that the attention is under your control, or being deliberately guided. In the last case—of sudden bright flashes and loud bangs—the way attention is grabbed is not under your control. In these cases attention looks more like an automatic reflex. Interestingly, substantial empirical evidence suggests that a distinction in the kinds of cognitive and neurological processes which underlie the two forms of attention (Carrasco 2011, 1488–9). Roughly, all these processes are modulations of the response rates in neurons within areas responsible for perceptual processing (Carrasco 2011, 1485). The responses of neurons encoding information about attended things are stronger than those encoding information about nonattended things. But this modulation can come from either lateral inhibitory signals between neurons (Desimone and Duncan 1995; see Ruff 2011; Carrasco 2011, 1486 for overview), or from back projections originating in higher areas of processing. In the first case neural responses to strong stimuli beat out or inhibit responses to weaker stimuli, suggesting (albeit tentatively) that this modulation underlies involuntary attention. In the second case the back projections largely originate in a single place, the frontal eye fields (Armstrong 2011). As the frontal eye fields

---

[10]The two options might come to the same thing, depending on the precise relationship between the personal-level action of attending and the underlying subpersonal processes. Some subpersonal process might be wholly constitutive of the personal-level action, collapsing the two views.

seem to be mainly involved in controlling voluntary shifts in attention, a natural (albeit very tentative) suggestion is that modulation from them underlies voluntary attention (see Carrasco 2011, 1489 for a similar suggestion).

Although Campbell never clearly articulates the perceptual link problem, his claim that it's attention which provides "knowledge of the reference" of demonstratives seems to be roughly the claim that attention is the exploited perceptual link (Campbell 2002, 22; see also Clark 2006, 168). In any case, Campbell's work is clearly the beginning of attention-based views on the perceptual-link problem, and is well known, so it's worth making two remarks. First, it's unclear whether Campbell intends for the restriction made here to voluntary attention. He refers to the relevant kind of attention as conscious, but isn't clear on whether by 'conscious' he means attention to consciously perceived objects, voluntary attention, or personal-level attending (vs subpersonal attentional processes). Second, Campbell assumes what he calls the classical view (2002, 24), which is something like an acquaintance condition which says that what provides knowledge of reference of mental demonstratives (what provides the perceptual link) must be what "causes, and justifies, your use of particular ways of verifying and finding the implications of" thoughts involving those demonstratives. Much of his account is an explanation of how attention acquaints the thinker with the target in this way. My project here, to reiterate, cleanly separates the perceptual link problem from explanations of how a purported perceptual link could provide the needed (if any) acquaintance; so details about how voluntary focal attention might do so are infelicitous.

3.3.2   Motivations

Why think that voluntary focal attention is the exploited link? First, the reason is because when you make intuitive assessments of whether a given case (e.g., example 15 vs. 16) involves perception-based thought you consider whether voluntary attention fixes the target. Second, when making intuitive assessment of the identity of a target you rely on what's being voluntarily attended. More generally, it's how voluntary attention is used when making intuitive assessments about cases of perception-based thought which provides preliminary motivation for the view. Before expanding these points, consider intuitive assessments and the methodology available in investigating perception-based thought.

The perceptual link problem is empirical: perception-based thought is a phenomena specified ostensively through example. What link is exploited is something to be discovered by investigating cases. Unfortunately, there are no tests

using observable behavior—even very indirect ones—for when a case of thinking is perception based. There also aren't any tests using observable behavior for deciding the target in cases of perception-based thought. A good contrasting example in which these tests are available is attention. Attention typically results in observable, or behaviorally detectable, effects, e.g. experiential changes (accessible through verbal reports) and improved sensory processing (accessible through reaction time, Posner 1980), and whether there is attention to a given thing in a certain case can tested by looking for these effects. There aren't any obvious or widely accepted observable effects in the case of perception-based thought.

This leaves two ways of investigating perception-based thought. The second is to use known facts about human or animal psychology to rule out, constrain, or suggest possible answers to the perceptual link problem. This approach is used in §3.4 to argue against more permissive attention-based views. The first is intuitive assessment of cases, which I use here to motivate the voluntary focal attention view and later in §3.5 to argue against the dominant causal source view. In intuitive assessment a case is presented to you (e.g., example 15) and you give an intuitive assessment of whether it involves perception-based thought and, if so, of the target. When put explicitly the intuitive-assessment approach can seem inappropriate as a methodology for answering an empirical question, but it should be pointed out that it's the standard starting point for most empirical investigations. For example, before attention could be operationalized using its behaviorally detectable effects as a test there had to be intuitive assessment of cases which counted as attention from which behaviorally detectable effects could be gleaned. The difference with perception-based thought is that there's no obvious or widely accepted operationalization which allows for testing that goes beyond intuitive assessment.

Returning to motivations, imagine you're making an intuitive assessment of a case like example 15 or 16. In example 15, but not 16, it's natural to take the target to depend on your voluntary attention: when you introspect your judgment, what matters for determining the target is your voluntary attention. Similarly, when you attribute perception-based thoughts to others it's natural to assign the thing to which they're attending as the target. In contrast, things not voluntarily attended aren't even potential candidate targets. Imagine that in example 15 there's a baseball next to the tennis ball; there's no way to run the situation so that you're voluntarily focusing attention on the tennis ball but still, somehow, have a perception-based thought about the baseball. Of course, you could think a descriptive thought which picked out the baseball via a condition (e.g., *the base-*

ball), even one that leveraged perception (e.g., *the seen ball next to what you're attending*), but this wouldn't be directly exploiting your perceptual link.

That intuitive assessments use voluntary attention in this way is often taken for granted (Raftopoulos and Müller 2006, 253; Dickie 2010, 234; Jeshion 2010a, 1; Levine 2010, 178). For example, perception-based thought is often introduced with descriptions of cases that specifically refer to the role of attention in selecting the target. More generally, it's natural to talk about perception-based thought in terms of attention. Campbell (1997; 2002) was the first to explicitly point out how attention plays into perception-based thought. He pointed out that attention is necessary for it (see also Levine 2010, 179; Dickie 2011, 298; Smithies 2011a, 265). To adapt and simplify an example (Campbell 2002, 8), imagine you're sitting at a dinner table in a crowded restaurant. You look around the room in a way so that you don't attend to any one person or thing; your visual experience is a "sea of faces". Having a perception-based thought about a person in view requires that you, at least for a moment, stop gazing around and voluntarily focus attention on one person. Important for showing that attention is what's missing, adding nonattentional information channels doesn't seem to help. Even if, similar to a case of super blindsight (Block 1995), you suddenly find yourself with a disposition to make unprompted judgments on the basis of visual information, and those judgments reliably track features of one person in the room, these still (the intuition goes) would not be thoughts in which the target was fixed by a perceptual link.

3.4   The Permissive Attention View

The first alternative to the voluntary focal attention view is one that identifies the perceptual link with attention more broadly, allowing that both voluntary and involuntary attention can fix the target of thought. On the permissive view the exploited perceptual link is the relation that holds between you and a consciously perceived thing in virtue of your attending to it, regardless of whether it's voluntary or involuntary. An example initially suggests this permissive attention view.

> **Example 17:**
> Consider again example 15 and imagine that while you stare at the tennis ball on the table a second tennis ball is thrown past you. As the ball flies past you the sudden motion in your field of view draws your attention to it. So, you're involuntarily attending to the flying ball.

Intuitive assessment of the case suggests that perception-based thought is possible: your involuntary attention to the ball affords a way to select it. You could, e.g., judge that what had grabbed your attention was another tennis ball in a way that would naturally be expressed by saying "*that* is a tennis ball".

In this section I argue that, despite what this example initially suggests, the more permissive attention view faces a fatal underdetermination problem. Empirical research suggests that there's a mechanism, attentional spreading, which when attention is initially directed at one thing (voluntarily or involuntarily) leads it to be involuntarily grabbed by other things in your perceptual field. Since this mechanism is in play generally, there won't be a problem with disagreements over intuitive assessments. If correct, then even in paradigmatic cases like example 15, which everyone accepts allow for perception-based thought, there is involuntary attention to consciously perceived things besides the selected target. In example 15 voluntary attention will be focused on the tennis ball, but attentional spreading will lead to involuntary attention to other items in the scene. So, identifying the perceptual link as attending (whether voluntarily or involuntarily) would make it underdetermine the target in a case (example 15) in which, intuitively, conscious perception provides an exploitable link. Thus, this basic fact about human psychology—that the mechanism of attentional spreading leads generally to divided involuntary attention—rules out the attention relation, more broadly construed, as the exploited link.

Before discussing attentional spreading, consider again example 17. Despite first appearances, it's not a case in which involuntary attention is exploited to select a target. When intuitive assessment suggests that perception-based thought is possible, what's being imagined is a more extended version of the case in which after the flying tennis ball grabs your attention you voluntarily "lock on" and hold your attention on it. Then, with attention voluntarily focused, you being thinking about the ball in a perception-based way. But once voluntary attention is engaged it's no longer a case of exploiting involuntary attention, despite the fact that it initially started out as an involuntary case. In contrast, when example 17 is run so that all you have is the attentional grab, with no subsequent voluntary holding of attention, my intuitive assessment is that perception-based thought is not possible. Of course, memory-based thought and descriptive thought (e.g., *the thing that just grabbed my attention*) are possible, but these aren't the same as perception-based thought.

3.4.1 Attentional Spreading

I begin by providing an overview of the two empirical results which suggest that there's attentional spreading: fMRI data on the "spread" of attentional modulation, and behavioral data on reaction time. It is well known that locations in some areas of visual cortices map topographically to positions in the retinal image (see Wandell and Winawer 2011 for review). The basic result relevant here is that if you compare the level of cortical activation in fMRI studies between a control case in which the subject attends to a fixation point and a case in which the subject covertly attends to a spot away from the fixation point, the level of difference (the attentional modulation) is most in the cortical location that corresponds to the attended spot. But, there is also a significant difference in activation (suggesting some level of attention) in the areas around that location.

In one study of modulation effects in areas V1/V2, Ritobrato Datta and Edgar DeYoe (2009) showed that the area of attentional modulation can be quite large and is not distributed equally around the attended spot. They also found that the pattern of modulation depends partly (at least in their experimental setup) on the position of the covertly attended spot relative to the fixation point. To show that the attentional modulations outside the cortical location corresponding to the attended spot are behaviorally relevant Datta and DeYoe ran trails in which they asked the subject to covertly attend to a spot of their choice and then, using both the location of maximum recorded enhancement and the overall pattern of modulation, predicted (without prior knowledge) the attended spot (out of 18 choices). When the maximum recorded modulation was used Datta and DeYoe were 100% accurate, while they were 87% accurate when using the overall pattern (2009, 1043). Overall these results suggest that when one location is attended the area around it draws attention as well, leaving attention divided between multiple locations.

Behavioral tasks extend these neurobiological results (see Scholl 2001, 8–9 for review; see also Dickie 2010, 216). In one type of task (based on Posner's cuing task, Posner 1980) a subject fixates on a cross in the center of a screen. A cue of some kind then flashes somewhere on the screen and grabs the subject's attention. Shortly after the cue the subject is given a discrimination task of some type, e.g. a letter appears on the screen and the subject must push a button corresponding to it. Typically discriminations are faster when the letter is at the location of the cue, but the effects extend beyond just that spot. Tasks like this suggest that cues can draw attention not just to spatial locations, but also to fea-

tures (like color and shape) and objects. For example, a cue appearing in a red box on one side of the fixation point will speed reaction time to letters appearing in a red box on the other side of the fixation point, suggesting that the cue has drawn attention to the color red, and not just the location of the cue. Similarly, a cue appearing at the top of a rectangle will speed reaction time to letters appearing at the bottom of the rectangle, suggesting that the cue has drawn attention to the rectangular object itself. In a recent set of experiments, Dwight Kravitz and Marlene Behrmann (2011) show that all three of these cuing effects (space-, feature-, and object-based) happen together, suggesting that when attention selects one instance of a location, feature, and object it also spreads to select the others.

The crux is that both sets of results suggest that involuntary attention is always divided. Both suggest that when you attend, either voluntarily or involuntarily, to one thing and your attention spreads to other things, such as nearby locations, the whole cued object, other instances of the cued feature. But, of course, your attending to these other things is not voluntary. It might be objected that spreading effects are a feature of subpersonal attentional mechanisms, not an instance of involuntary attention being divided. What spreads is lateral inhibitory signals between neurons, not necessarily the focus of personal-level involuntary attention. But, there is no reason to block the move from attentional spreading to personal-level acts of attention and at least some reason to make it. For example, attentional spreading has the behavioral effects of attention (e.g., decreased reaction time). It's certainly not implausible that, thanks to the spread of attentional enhancement, when you attend to a certain location you also end up attending to the adjacent locations, or that when you attend to an object that you also end up attending to its location and further to adjacent locations.

3.4.2 Descriptive Mediation

With the evidence for attentional spreading summarized, I turn to one way the permissive view might be defended. The objection I've given assumes that a relation which would leave the target underdetermined cannot be exploited to select a target. You might suggest that this assumption is only correct if the relation has to do all the work in fixing the target. But the means of selection used in perception-based thought, the reply continues, might not only rely on the perceptual link to fix the target. Perhaps selection in perception-based thought involves using some descriptive resources in addition to the perceptual link, similar to how complex demonstratives can be used in speech acts. Imagine, for example,

that you point to a chair with a few books on it. Even if your pointing—or whatever demonstrative mechanism is used in demonstrative speech acts—doesn't discriminate the chair from the books, you can still demonstratively refer to the chair by saying "that chair". The reply in defense of the permissive view is that maybe selection in perception-based thought works in this way. (The reply is partly a rejection of the presupposition that perception-based thought is demonstrative in the first sense from §3.2.1, i.e. is direct.)

The reason this reply fails is that it only works if selection in perception-based thought is always descriptively mediated, but this is highly implausible.[11] I think it's quite plausible that many cases of perception-based thought involve just this kind of descriptive mediation. Introspectively, I often rely on both my conscious perception of an object and my grasp of concepts of some of its perceived properties to select it for thought, just as I often use both demonstration and descriptive terms in complex demonstrative speech acts. But there are many cases, like example 15, in which intuitive assessment of the case suggests that no descriptive resources are being used, along with a perceptual link, to select a target. Since attentional spreading is a mechanism that produces division in attention in all cases, it happens in these cases too. So, the appeal to descriptive mediation won't be available.

3.4.3 Divided Voluntary Attention

Finally, I conclude this section by considering whether the voluntary focal attention view faces its own underdetermination problem, albeit a different one. Section 3.3.1 said that the restriction to *focal* voluntary attention was required because when voluntary attention is divided it leaves the target of thought underdetermined. But you might object that perception-based thought is possible in cases when voluntary attention is divided, and so that any answer to the perceptual link problem must accommodate them. In that case, underdetermination can't be avoided by a stipulated restriction to focal voluntary attention. What's left—an account that identifies the link as voluntary attention—will face an underdetermination problem in those cases in which voluntary attention is divided.

How would you decide whether perception-based thought is possible in cases when voluntary attention is divided? It comes down to intuitive assessment:

---

[11]Keep in mind that the discussion is between those who accept that there is a problem about perceptual links. Of course, a descriptivist about perception-based thought (see §3.2.1) who flatly rejects its minimal directness will say that all perception-based thought is descriptively mediated, but won't have been on board with the perceptual link problem from the start.

cases need to be described, from a first or third-person point of view, then an intuitive verdict returned on whether in those cases it's possible to select directly a consciously perceived thing. For example, consider the case mentioned above of voluntarily splitting your attention between two adjacent words on the page you're now reading.[12] When you voluntarily attend to two adjacent words, can you have a perception-based thought about one of them without first breaking attention off the other and focusing attention on just that one? My intuitive assessment is that I cannot.

This intuitive assessment about the possibility of perception-based thought in cases of divided voluntary attention has been disputed. For example, although he doesn't elaborate, Levine says that he "can attend to two objects while explicitly thinking [that one] of one of them" (2010, 179). He concludes that a purely attention-based view cannot be correct and that "a separate mental act of demonstration is necessary" for perception-based thought (2010, 178). The following case helps motivate Levine's claim (personal communication, Indrek Reiland).

Example 18:

You voluntarily hold attention on two adjacent words. Then, without breaking attention from one or switching attention between them, you think a perception-based thought expressed by saying "*that* [one of the words] is shorter than *that* [the other word]". Here you sequentially exploit conscious perception to select one of the words and then the other, during a thought, while keeping attention split

---

[12]Pylyshyn's work on multiple object tracking (MOT) (see Pylyshyn 2007) is so well cited, including in connection with divided attention and perception-based thought (e.g. Levine 2010, 178), that it's impossible not to situate it within this discussion. I assume the reader is familiar with the work, but the gist is that subjects are asked to track a number of flagged objects (around 4) on a screen with (say) 8 objects total. It turns out, rather surprisingly, that people can track up to 4 objects rather well. Pylyshyn uses details of this result, not mentioned here, to infer facts about the cognitive processing underlying vision. Specifically, he postulate what he calls a FINST mechanism as a tag to keep track of objects in early visual processing. It's quite plausible that voluntary attention is divided (and used) in MOT tasks (see Scholl 2009, 55). But for the discussion here, the MOT tasks are no different from the every day examples like splitting your attention between two adjacent words on the page. My response to these cases will apply straightforwardly to the MOT cases: voluntary attention is divided, but intuitive assessment suggests that perception-based thought isn't possible. The MOT cases are only interesting if there's something unique about them which suggests the intuitive assessments should come out differently, and I don't see any such features. Finally, note that Pylyshyn actually denies that MOT involve divided attention, or at least that attention is doing the tracking. He says the tracking is done by the FINSTs, and these are distinct from attention (Pylyshyn 2001, 147,152, 2007, 39,44,60,94–96; Scholl 2009, 53; Dickie 2011, 301–2).

> between them (i.e., without sequentially focusing attention on the words).

My intuitive assessment is that this case is impossible. When I try to think a thought like in example 18, I don't feel I can do it. Sequentially focusing attention on the two words seems necessary to me.

Someone who shares my intuitive assessment might reply in three different ways. First, she could suggest that example 18 may not require completely breaking attention from one word, but does require a sequential shift in the emphasis of attention. For the first selection your attention is primarily focused on one word, then for the second you shift the primary focus of attention to the other. (Then the voluntary focal attention view only needs to be further refined, so that the target is the primary focus of voluntary attention.) Second, she could suggest that those who think example 18 is possible are really using descriptive means to resolve the underdetermination. Third, she could ask *which* one of the attended words has been selected first and which second. Those who say example 18 is possible cannot appeal to their attentional state to answer this question, since (they say) it remains equally divided. Following Levine, they might appeal to some extra-attentional selection mechanism (what Levine calls mental pointing). But even if there is an extra-attentional selection mechanism, it's not a cognitive process to which you have conscious access. While I feel familiar with my own attending, introspection reveals no selective pointing process beyond it. For example, I wouldn't know how to follow instructions to hold attention steady while shifting my mental pointing. The upshot is that it's unclear what basis someone could have for saying example 18 is possible.

3.5   The Dominant Causal Source View

Starting with Evans (Dickie 2010, 227), a widely received answer to the perceptual link problem identifies the perceptual link as the relation which holds between you and a consciously perceived object in virtue of it being the dominant causal source of information (Evans 1973, 1982; Raftopoulos and Müller 2006, 253). Typically the view is tied to the mental file framework, so that the target of thought is the dominant causal source of information in some mental file. While §3.4 used known facts about human psychology to rule out the permissive attention view, this section will rely on intuitive assessments to argue against this dominant causal source view. Specifically, I give two examples and argue that in each case the dominant causal source view disagrees with what you'd intuitively

assess as the target.

Before giving the examples, it's helpful to summarize the rudiments of mental files. The term 'mental file' covers a variety of different (but related) theoretical constructs (see Recanati 2010, 156 for references), but the basic idea is that we or our underlying cognitive systems can usefully be thought of as keeping a "file" of information on each thing with which we're acquainted (Bach 1987; Raftopoulos and Müller 2006, 256–58,264; Pylyshyn 2007, 37; Dickie 2010, 222–25, 2011, 304; Jeshion 2010c, 129–35; Recanati 2010, 156, 2012, 34–38). Roughly, a mental file is functionally defined as a collection of information or beliefs which our cognitive processes treat as being about the same thing. The referent of a particular file, the thing about which it's collecting information, is the dominant causal source of information in the file.

How do mental files fit into an account of perception-based thought? Although explicit details are rarely given, the usual approach is in terms of mental demonstratives (see §3.2.1). The idea is that certain mental files, or a single one, are mental demonstratives (Bach 1987; Raftopoulos and Müller 2006, 256–8,264; Dickie 2010, 222, 2011, 305; Jeshion 2010c, 135; Recanati 2010, 157, 2012, 34–8, 57–67). Presumably these files are special-purpose temporary ones with the functional role of storing information about what's currently being consciously perceived. They are opened to store incoming perceptual information and closed (and reopened) as what is being perceived changes, or is taken to change (Jeshion 2010c, 131; Recanati 2010, 157, 2012, 35–7,62). For example, when you see the tennis ball in example 15 a file is opened and information about the ball gained through you conscious visual experience of it is added to that file. Tokening a mental demonstrative in a perception-based thought amounts to using or deploying these special-purpose files. The target of a given instance of perception-based thought is the referent of the mental demonstrative tokened in it, which is just the referent of the particular instance of this special mental file used in the thought (Dickie 2010, 222).

3.5.1 Case 1: Synchronic Multimodal Perception

For the first example, imagine that you see a tennis ball rolling across a table and hear a noise, a consistent sound you mistake as the sound of the ball rolling. The mistake isn't just at the level of deliberate judgment, say, but is a perceptual binding error. You mistakenly hear the sound as coming from the ball, in just the same way as you mistakenly hear a ventriloquist's voice as coming from her doll. Perhaps the tennis ball makes no audible sound as it rolls, and what you hear is the

sound of a steel bearing rolling across a table behind you. Now in this situation you voluntarily attend to just the noise, while the ball stays in view. Since the mistake is a binding error, there's presumably a single mental file collecting information both from the seen ball and the heard noise. Although it's unclear how to quantify amounts of information, it's not implausible that in this case much more visual information is being collected than auditory information. Hence, the tennis ball is the dominant causal source.

Now say you have a perception-based, e.g. you judge that the thing to which you have a perceptual link is a tennis ball. Intuitive assessment of the case suggests that the attended sound, or its source, is the target of thought. This assessment can be motivated in part by considering how you would evaluate the judgment once you learn that the sound isn't coming from the ball. In the case described, with my attention voluntarily focused on the sound, once I learn of the binding error I'd be inclined to say that my judgment was incorrect: *that* (the target of my thought) wasn't a tennis ball, it was whatever was making the noise from behind me. Or, consider ventriloquists again. Imagine that you're watching one for the first time, intently focusing on the sound of the doll's voice. Your eyes are still on the doll, but your attention is entirely, or at least primarily, focused on the voice. You think a thought you'd express by saying "I'd swear *that's* coming from the doll". Surely in this case the target of your thought is the voice to which you're attending, despite the fact that (again, because of the binding error) the mental file on the voice is also collecting visual information from the doll—and plausibly much more of it.

3.5.2   Case 2: Diachronic Unimodal Swap Case

In the second example, imagine again that you're watching a tennis ball roll across a table. This time there's a block on the table which occludes your view shortly before the ball reaches the end. The result is that you watch the ball—voluntarily hold attention on it—for 10 seconds as it rolls across the table, for 8 seconds before it rolls behind the block, then just under 2 seconds as it rolls the rest of the way. The momentary occlusion of the ball doesn't cause a new mental file to be opened; instead you see it was a continuously rolling ball. But the setup is a trick: behind the block the ball is swapped with an indiscriminable one. The ball you tracked for the first 8 seconds is not the ball you tracked for the last 2 seconds.

Imagine that as the ball reaches the end of the table—at the end of the 10 seconds—you make some kind of perception-based judgment about it. Based on intuitive assessment, what is the target of the thought? There are three possible

answers: the first ball you saw for 8 seconds, the second you saw for 2 seconds, or that there's reference failure of some kind. My assessment is that the target is ball 2 (the one you saw for 2 seconds at the end). I can see the suggestion that there's a reference failure, but I think there's no intuitive plausibility to saying that the target is ball 1. Just imagine: you watch the events just described and at the end, while holding attention on ball 2 as it reaches the end of the table, think a thought you'd express by saying "*that* [looking right at ball 2] is . . .". It doesn't seem, as a matter of intuitive assessment, that you could conclude that the target is ball 1 in this case.

But, because of your longer exposure to it, it's plausible that ball 1 is the dominant causal source of information. If so, then this is another case in which the dominant causal source view gets the target wrong. The dominant causal source theorist might reply by saying that it's only ball 1 on an overly simplistic version of the dominant causal source view. This version has it that the target is the *overall* dominant causal source of information. Perhaps instead the target is the *current* dominant causal source of information—the dominant source at the time of thinking. Since ball 2 is presumably the dominant causal source of information at the time of thinking, this version would return the correct answer. The main problem with this response is that it would mean mental files can shift referents (as the current dominant causal source of information changes), but mental files cannot shift referents. After all, the point is that mental files collect information *about the same thing.* Of course, the dominant causal source theorist could offer other potential amendments (e.g., making a file's referent the *initial* dominant causal source of information), but there's no obvious ones which don't involve referent shifts but get the target right in this case.

3.6   Conclusion

I've tried to do the following in this paper. First, I articulated the perceptual link problem in a way that doesn't presuppose that perception-based thought is relational or demonstrative, and that frees it from an information-channel framework. Second, I've presented motivation for accepting the voluntary focal attention view as an answer to the problem. I then introduced the mechanism of attentional spreading and argued that it provides a constraint about human psychology which rules out more permissive attention-based views. Finally, I offered two examples in which the main rival to the voluntary focal attention view, the dominant causal source view, returns a target different from what you'd intuitively assess as the target.

Chapter 4:
The Role of Visual Representations in Seeing

4.1   Introduction

Seeing is an occurrent interaction between an organism doing the seeing and the thing seen. Although what makes an interaction an instance of seeing is a joint empirical-philosophical question for which there's no widely accepted answer, some claims are a pretty safe bet (e.g., Prinz 2006, 436). For example, seeing typically (necessarily?) involves a causal chain in which light reflects or is emitted from what's seen into the organism's eyes, thereby exciting photoreceptors in the retina. Next, it also typically (but not necessarily?) involves the organism having an experience of what's seen. Third, seeing involves processing information from the light hitting the retina. This processing involves the construction of representations of the distal causal sources of the proximal retinal stimulation. Given the success cognitive psychology and cognitive neuroscience have had approaching seeing in these information processing terms, a tempting—and widely held[1]—view is that seeing *constitutively involves* your visual system constructing a representation of what's seen. Seeing *just is* the construction of a representation by your visual system.

In this paper I argue that this view is false. Specifically, I give an example of a thing you see but for which no representation is constructed in the visual system. The example is seeing certain arbitrary parts of the figures used in a particular

---

[1]I'm unaware of any philosophers who have explicitly articulated and defended the view, but Jesse Prinz (2000, 249, 2006, 454, 2011, 174) and Michael Tye (1995, 100-03,120–23) seem to hold it. In psychology the view seems to be often tacitly assumed. For example, in a recent survey article on gestalt psychology Johan Wagemans et al. say grouping principles "pervade virtually all perceptual experiences because they determine the objects and parts that people perceive in the environment" (Wagemans et al. 2012, 1180). As I discuss below (§4.2.3), grouping principles describe the conditions under which the visual system constructs representations. So the claim is that these visual representations determine what's seen. But the view isn't always assumed. For example, in his work on the neural correlates of consciousness Logothetis (1998) leaves the question open.

71

multiple object tracking (MOT) task. This MOT task is from a study by Brian Scholl, Zenon Pylyshyn, and Jacob Feldman (2001). I show that they're seen but not represented in three steps.

Step 1: Using the results from Scholl et al. from their MOT study, I argue that during the MOT task itself the visual system does not construct a representation of certain parts of the tracked figures.

Step 2: Using introspection, I argue that during the MOT task these nonrepresented parts are available for voluntary attention.

Step 3: By appealing to a standard view of voluntary attention, I argue that because they are available for voluntary attention these nonrepresented parts are seen.

Note that the argument integrates three different approaches: step 1 uses empirical results, step 2 uses introspection, while step 3 appeals to broader theoretical considerations. Although the actual example used is seeing certain arbitrary parts of depicted figures while doing a MOT task, if the argument works it's plausible that the conclusion generalizes. The visual system does not construct representations of most arbitrary parts of seen objects, in most cases.

Before proceeding, note the following five preliminary points. First, the thesis—that seeing does not constitutively involve the visual system constructing a representation of the thing seen—is consistent with seeing constitutively involving the visual system constructing *some* representations. For example, seeing a part X of an object might involve constructing an overall representation of that object without constructing a separate representation of X. What's at issue isn't whether seeing constitutively involves the construction of representations by the visual system. I take for granted that seeing does involve these (as I'll sometimes say for short) visual representations.[2] Instead, what's at issue is the relationship between visual representations and the overall organism-level interaction of seeing itself. If successful, the argument here shows there isn't a one-to-one

---

[2]As I explain below, although some philosophers (most naïve realists) deny that seeing is a "representational state", this is compatible with them accepting—as most do—that it constitutively involves the construction of representations by the visual system. But this assumption *is* rejected by ecological psychologists, enactivists, and some working in AI (e.g., Gibson 1966, 1986; Brooks 1991; van Gelder 1995; Noë 2004; Hutto and Myin 2013; see also Orlandi 2011a,b, 2014). Kathleen Akins' work is of note. She (1996, 364–68) gives a careful neurologically grounded critique of the view that sensory systems construct representations *while also* (as I do here) articulating the gap this creates between the intentionally directed character of perceiving and the underlying sensory systems which enable perception.

mapping between things seen and visual representations. Instead, the visual system engages in a fairly limited and circumscribed construction of representations which underlies or enables seeing a rich array of objects and their parts, features, properties, locations, and spatial relations.

Second, it's a metaphysical or ontological question whether seeing a thing constitutively involves the construction of a visual representation of it. It's not a question of causation or explanation. Water and lightning provide familiar examples. The discovery that lightning is an electrical discharge was a discovery about constitution: it's not that discharges of electricity cause or explain lightning, instead they're what it's "made" or consists of. Likewise, that water is H_2O is a metaphysical or ontological fact about its constitution. Just as we know lighting consists of electrical discharge and water consists of H_2O, you might suggest that seeing constitutively involves constructing a visual representation of what's seen. Constructing a visual representation of a thing based on retinal stimulation is just what it is to see that thing, the suggestion goes. Giving an example of seeing without a visual representation of what's seen will show that the suggestion is false, just as a case of lightning without electrical discharge would have shown that electrical discharge isn't constitutive of lightning.

Third, the thesis here is specifically about *conscious* seeing: seeing in which the perceiver has a visual experience of what's seen. So the thesis is that the construction of a visual representation of a thing is not constitutive of having a visual experience of it.[3] What's meant by 'conscious' or 'experience'? Seeing is conscious, or involves an experience of what's seen, when there is some subjective, first-person perspective had by the perceiving organism. To use Thomas Nagel's phrase (Nagel 1974), seeing is conscious when there's "something it's like" for the organism seeing. To use another suggestive way of speaking, seeing is conscious, or involves an experience of what's seen, when there's some way what's seen *looks* or *appears* to the organism. For example, what's in the blind field of blindsighters is nonconsciously seen because there is no way it looks or appears to them. Also note that consciousness, in this sense of having an experience, shouldn't be confused with availability of information, or what Ned Block (1995) calls access consciousness. As Block points out, you can imagine a blindsighter regaining spontaneous access to visual information for use in reasoning and verbal report without regaining experience, i.e. without what's in their blind field having a look or appearance. In this paper I use the two terms 'conscious seeing' and 'experience'

---

[3]When I refer to experiences I always mean nonhallucinatory ones, i.e. kind you have when you successfully and consciously see.

interchangeably, selecting whichever leads to the most lucid wording.

Fourth, the question posed by this paper (about the relation between visual representations and what you visually experience) is orthogonal to the central debate in philosophy of perception between representationalism and naïve realism. *Representationalist* views hold that experiences are representational states (Tye 1995; Dretske 2003; Burge 2005, 2010; Pautz 2009; Speaks 2009; Siegel 2010a; Schellenberg 2011). They're representations of what's experienced and thereby have content which can be assessed for veridical or accuracy. Any two instances of experience, on this view, will be instances of the same experience—will be the same experience—iff they have the same content. *Naïve realist* views, on the other hand, hold that experiences are relational states (e.g., Campbell 2002; Martin 2004; Brewer 2011; Johnston 2011; Logue 2012b; Hobson 2013; Knight 2013). They're relations to what's experienced. Any two instances of experience, on this view, will be the same experience iff they involve being in the experience relation to the same thing(s).

Denying that seeing a thing constitutively involves visual representations of it is consistent with representationalism. Few representationalists would speculate *a priori* about the relation between experiences and visual representations. Almost none assume that having a visual experience of a thing constitutively involves a visual representation of it (see fn 1). In addition, both the assumption that consciously seeing a thing constitutively involves some visual representation and the further (here rejected) claim that it involves a visual representation of the thing experienced are consistent with denying that experiences are representations and accepting that they're relations. Both points are explained by noting that the nature or character of experiences—representational or relational—need not be "directly grounded" in what happens in the visual system (even if what happens in the visual system is constitutive of experiences).

Finally, Zenon Pylyshyn has also argued that a visual representation of a thing is not constitutive of having a visual experience of it (2007, 120–23).[4] His argument starts with his FINST account of MOT results. (I explain FINSTs below; for now it's enough to know that FINSTs are a computational mechanism Pylyshyn posits to explain results from MOT experiments.) As he presents it, the FINST

---

[4]Others have also done important related work (see Chalmers 2000, 87; Noë 2004, 37; Noë and Thompson 2004; see also Dennett 1978; McDowell 1994a). For example empirical example, Stephen Mitroff et al. (2005) have done work on the stream-bounce ambiguous display suggesting (e.g.) that the visual system can treat the targets in the display as bouncing while the subject reports the experience of them streaming.

account entails that at any one time only a few objects are represented by the visual system (those being tracked by a FINST). But, he suggests, introspectively it seems that you have a panoramic visual experience of many objects. So, there must be some disconnect between what's represented in the visual system and what's consciously seen, i.e. the things of which we have visual experience.[5] Note that Pylyshyn's argument has the same structure as mine: first use empirical results to show the things for which there are visual representations, then argue that you have visual experience of more things.

My argument makes several advances. First, although step 1 also draws on MOT work, I don't assume Pylyshyn's FINST-based account of those results. Second, I use the empirical results to point to a specific example of something not visually represented: specific parts of a figure during a particular MOT task. Third, my approach to arguing that these parts are visually experienced does not rest on the folk panorama view of visual experience. Steps 2 and 3 of my argument are a sophisticated attempt to show that these parts are visually experienced without appealing to the panorama view, or any intuitive introspective judgments about visual experience at all. I take this lack of direct appeal to introspection of visual experiences, along the integrative approach, to be a strength of my argument.

4.2   Step 1: Visual System Representations

There are many things you see. A plausible list includes (1) lighting (e.g., a flashlight pointed at a dark wall, the effects from a colored bulb), (2) shadows, (3) temporally extended events (e.g., a ball rolling), (4) spatially connected medium-sized objects (e.g., a chair, other people), (5) large static objects that extend beyond your field of view (e.g., a nearby building or mountain), and (6) contiguous masses (e.g., water in a puddle). In addition, you also see the (7) surfaces, (8) parts (e.g., the leg of a chair), (9) low-level features (e.g., color, shape, luminance, texture), (10) high-level properties (e.g., being a chair, being a specific person), and (11) and spatial relations between objects and masses. Here the focus will be on objects and their parts, or at least on object-like figures depicted

---

[5]Pylyshyn suggests a number of arguments (see 2007, ch4,5), but the one just outlined is the clearest. His discussion is also much broader than mine, hitting on issues related to visual imagination as well. He's also concerned with arguing against the view that visual representations are pictorial and that the retinotopic maps in V1 realize these representations (2007, 139–43). My aim here isn't to argue against any specific views about the form of visual representations and my argument doesn't presuppose any answer to this question.

in MOT tasks and their parts. Using a MOT task from Scholl, Pylyshyn, and Feldman (2001), I'll give an example of something not represented by the visual system.[6]

4.2.1   Preliminaries on Visual Representations

But what is the visual system and in what sense does it construct representations? Seeing is a task, something organisms do. Cognitive psychology gives an explanation of how they do it by breaking it into subtasks, each simpler than seeing (Drayson 2012). Each of these subtasks is functionally defined and carried out by, or realized in, physical neural processes in the brain. Often these subtasks are to compute some mathematical function, i.e. to compute some output as the function of some input. These functionally defined subtasks and the realizing neural processes are often called *subpersonal* tasks and processes to distinguish them from the personal-level activity of seeing itself. The *visual system* is the collection of neural processes realizing the functionally defined subtasks involved in seeing. The visual system is *constructive* in sense that the functions it computes take as inputs (for example) encodings of retinal stimulation and have as outputs encodings of the distal causal sources of that stimulation—of what's seen. What neural activity in the retina directly encodes is the distribution of light intensity on the retina. Seeing requires working backwards from encoded distributions of light intensity to their causes, i.e. to what's seen (e.g., Marr 1980, 203; see Fodor and Pylyshyn 1981 for discussion). The input and output states of these computations are *representational*: the input and output encodings can be more or less accurate and are about what's seen (Burge 2010, 292; see also Orlandi 2014, 9–15). The output encoding, e.g., might characterize what's seen incorrectly, attributing to it the wrong color or spatial location.

David Marr's work on shape extraction provides a well known example of this constructive process (see Marr 1980; see also Marr 1982).[7] It will also serve as a running example for explanatory purposes throughout this section (§4.2). Marr suggests that the visual system has (among many others) the subtask of taking an encoding of the light intensity distribution on the retina and outputting an encoding of the shape which produced that distribution, an encoding suitable for iden-

---

[6]Hereafter I sometimes refer to things for which the visual system constructs representations as being represented by the visual system.

[7]It's true that Marr's work is dated and the general explanation is almost certainly wrong. Still, it will serve as a relatively simple, well-known working example for more general points about perceptual processing and the construction of representations which are current.

tifying the shape (Marr and Nishihara 1978; Marr 1980, 211). The encoding of the retinal light intensity distribution can be thought of as an array with an intensity value at each point in the array (Marr 1980, 203). The output, which Marr calls a 3D model (1980, 211), encodes the overall shape by listing the simple component shapes that make it up and their relations to one another, given in object-centered coordinates. Marr divides this task into three further subtasks: (1) compute the light intensity changes in the initial input array (the "primal sketch"), (2) compute, from this primal sketch, the visible surfaces elements of the shape, including their depth and orientation, in viewer-centered coordinates (the 2½D sketch), (3) compute, from this 2½D sketch, the 3D model. A complete account of shape extraction would require specifying the algorithms the visual system uses to make these computations. Although Marr doesn't theorize about the neural realizers of this processing, one suggestion is that the primal sketch is computed in the primary visual cortex, area V1, some neurons of which are known to fire only in response to "edges" (changes in light intensity) with specific orientations (Hubel and Wiesel 1959, 1962, 1968; but see Chirimuuta and Gold 2009). Other evidence suggests that the 2½D sketch is computed in extrastriate cortex (V2–MT) and the 3D modal is computed in the inferior temporal (IT) cortex (Prinz 2000, 245). Whatever they are, the neural realizers of each step contain (respectively) representations of light intensity changes, surface elements of the shape, and of the full 3D shape. Alternatively put, the output state of the first step represents an array of light intensity changes, of the second step the visible surfaces of the shape, and of the final step the shape itself.

Importantly, the construction of representations by the visual system can be divided into two operations: feature extraction and grouping (Kahneman et al. 1992, 176–8; Scholl 2001, 16; Scholl et al. 2001, 160; Wagemans et al. 2012, 1180).[8] Feature extraction involves the detection of a feature (e.g., shape, color, luminance, or orientation) based on input from an earlier stage of processing (or directly from retinal stimulation). Marr's work, for example, provides an account of shape extraction. Grouping involves what's typically called *binding* clusters of those features together so as to treat them as belonging to the same thing (see figure 4.1). Grouping results in a *segmentation* of features in the scene into discrete chunks, or into distinct perceived things. On an influential account (Treisman and Gelade 1980; Kahneman et al. 1992; Treisman 1998), this grouping process is temporally extended and involves tracking a changing cluster of features over

---

[8]Austen Clark (2004) provides a longer discussion of many of these issues.

time in an "object file". The distinction between grouping and feature extraction is relative to a level of processing (Wagemans et al. 2012, 1188). For example, shape extraction, as described by Marr (1980, 211), itself involves grouping of previously extracted features (grouping edges when moving from the primal sketch to the 2½D sketch, and grouping shape parts when moving to the 3D model).

It's generally thought that there are a number of different mechanisms responsible for grouping, the mechanism(s) operating at a given time being task dependent (Kahneman et al. 1992, 178; Scholl 2001, 20–21,31; Scholl et al. 2001, 161; Wagemans et al. 2012, 1188,1205). Some of these mechanisms (e.g., Pylyshyn's FINSTs, which are discussed below in §4.2.2) are located within visual processing itself and are wholly stimulus driven. They don't draw on higher-level knowledge from cognitive systems and aren't under voluntary control. But there are also non-stimulus-driven grouping mechanisms which depend on input from cognitives systems. Attention is one example, but some grouping phenomena (e.g., seeing the dalmatian in figure 4.1) suggest that there are also nonattentional, non-stimulus-driven mechanisms as well.

The basic point is that the construction of a representation of a thing in the visual system happens either when it's a feature that's extracted by the visual system, or is a thing the features of which the visual system groups together. The output states of feature extraction and grouping operations represent (respectively) the extracted features and the things which have the grouped features. Note that, in the literature on grouping, often anything the features of which are grouped is *called* an object, or a "visual object". But so-called visual objects can include more than what I called objects in the above list of things seen (Kahneman et al. 1992, 178). The grouping of the edges of a visible surface, for example, would be a representation of that surface. Calling the things the features of which are grouped by the visual system "objects" doesn't limit the kinds of things that get represented via grouping. If you want to know whether the visual system constructs a representation of something seen, the question is still whether that thing is an extracted feature or has features the visual system groups together.

As indicated by figure 4.1, grouping of features by the visual system typically leads to a change in experience (but not always, see Wang et al. 2012). That is, the phenomenology, or look, of a seen thing is different when the visual system groups together seen features of it. It's natural to capture this difference with the phrase 'see as'. For example, before the dalmatian pops out in figure 4.1, you don't see it's collection of spots *as* a distinct object. Once the dalmatian pops out—once the visual system groups together its seen features—you do see it as

a distinct object. This effect provides an experience-level correlate to the visual construction of representations via grouping, as well as a source of defeasible introspection-based evidence for these constructions.

Figure 4.1: At first the above picture looks like an incoherent array of botches. But after a moment you should recognize a dalmatian in the scene, along with a tree and leaves. Note that once the dalmatian is recognized there's a shift in your experience; the dog "pops out" and is actually seen as a dog. The shift provides an example of grouping: initially when seen as a mere array the visual system does not group the patches together, but does once their recognized as a dog (leading to the experiential shift). Michael Bach reports on his website that the image is from Gregory (1970) (photo by Ronald James), but first published in *Life Magazine*, 2/19/1965, p. 120. See http://www.michaelbach.de/ot/cog_dalmatian/.

4.2.2   MOT and Arbitrary Object Parts

The MOT paradigm (Pylyshyn and Storm 1988; see Pylyshyn 2007 and Scholl 2001, 9–10 for overview) provides behaviorial, nonintrospective tests for whether the features of a seen thing are being grouped together, and so represented, by the visual system (Scholl et al. 2001, 161).[9] In MOT tasks a subject looks at a com-

---

[9]Other tests come from Posner's cuing (Posner 1980; Posner et al. 1980; see also Scholl et al. 2001, 8–9) and the preview-effect paradigm (Kahneman et al. 1992). To explain one, in a typical Posner-style set up a cue is given which draws attention. It's known that attention increases

puter screen which displays a number of "objects" (around eight), e.g. crosses, disks, or squares, none of which have any features that would allow them to be distinguished from the rest. Some of the objects (about four), called *targets*, flash or in some other way are cued. Then all the objects, including the uncued distractors, move in random, independent paths within the screen. When the objects stop the subject must pick out the target objects, for example by clicking on them with the mouse (Pylyshyn 2007, 34–35). The interesting result is that subjects can do this task at all, and with relative ease (Scholl 2009, 59). Further, performance remains relatively constant when there is up to four or five target objects, and then drops off significantly after that.

What's crucial for the discussion here is that the ability to track the targets indicates that the features of those targets (edges, shape, location, color) are being grouped together, and hence that the visual system is constructing a representations of the targets. So, a failure to successfully complete a given MOT task—to track multiple objects of a given type—indicates that the visual system is not constructing representations of the targets (Scholl 2001, 32; Scholl et al. 2001, 171–72). This is suggested by the two features just mentioned: tracking in MOT tasks is relatively easy and performance is constant to a point after which it falls flat. Both are characteristic of a fast, automatic perceptual operation like grouping (Dickie 2010, 220).

Note that much of the interest in MOT experiments concerns what they tell us about the involved grouping mechanisms. For example, a significant result is that performance in MOT tasks isn't helped when all the objects have different colors and shapes and isn't hurt when the objects randomly change colors and shapes during the trail (Pylyshyn 2007, 37). Pylyshyn uses this result to support this FINST theory, which says that the feature-grouping facilitating tracking in MOT tasks involves a preattentional mechanism which attaches labels, called FINSTs, to the feature groups. These FINSTs allow the visual system to keep track of targets without using a proxy like location or color (see Pylyshyn 2001, 2007; see Scholl 2009 for criticism).

A MOT tracking experiment from Scholl, Pylyshyn, and Feldman (2001) provides the example of a thing not represented by the visual system. In their exper-

reaction time and accuracy. The effect of the cue—which drew attention—on subsequent tasks is then tested. For example, the cue might be to flash or highlight one element of a figure in the scene. If increased reaction time or accuracy to changes in other elements of a figure is then observed, attention seems to also have been directed to them. Assuming that attention is drawn to whole perceptual groups, then would then show that those elements had been grouped together.

(a) Standard MOT Task

(b) Rubber Band Merging Condition
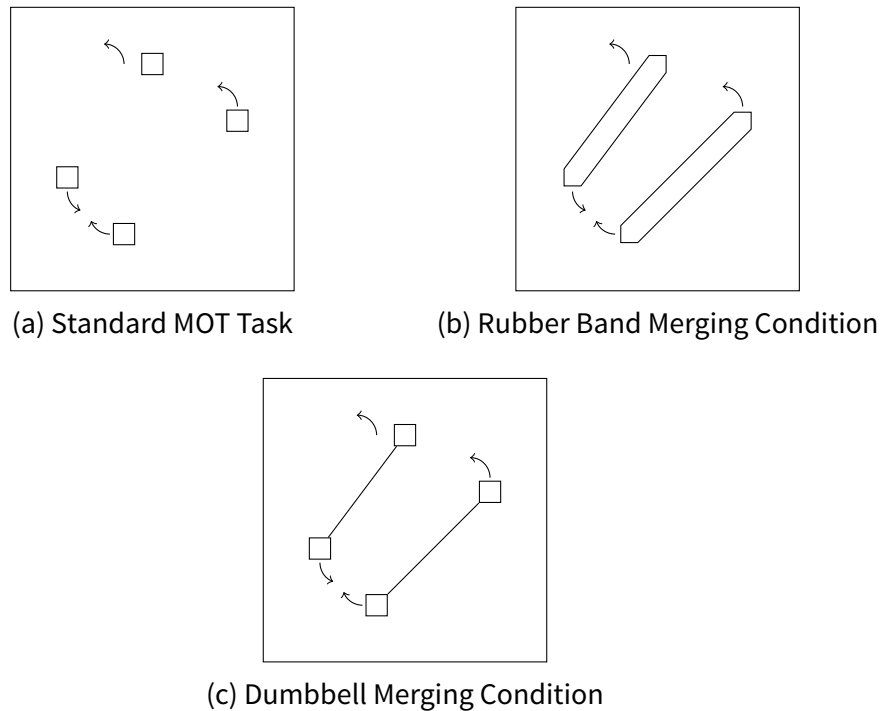

(c) Dumbbell Merging Condition

Figure 4.2: Schematic Diagrams of Scholl et al. MOT Task. (a) shows a standard MOT set up: several distinct figures (usually around 8, 4 displayed here) move at random around the screen. Some of these are cued at the start as targets (usually around 4, we can imagine the bottom two were cued here) and the subject tracks them as they move. (b) shows the rubber band merging condition. In this task the ends (both targets and distractors) move exactly how they did in (a), but are "merged" to form the elongated rectangles. (c) shows the dumb bell merging condition. Figures based on/adapted from (Scholl et al. 2001, fig.1 and fig.2).

iment they took the usual eight objects in a MOT task (eight boxes), paired each of the four targets with a distractor, then "merged" the paired targets and distractors (see figure 4.2). They used eight different merging conditions.[10] One of the merging conditions, the one which provides the example, wraps the boxes in a solid line, as if they've been wrapped in a rubber band (figure 4.2b). Another, mentioned here as a contrasting case, joins the boxes with a solid line so that the merged boxes form a dumbbell shape (figure 4.2c). The key result is that on some merging conditions, e.g. the rubber band condition, subjects are unable to track the targets, while on others, e.g. the dumbbells, tracking remains possi-

---

[10]See http://www.yale.edu/perception/Brian/demos/MOT-Merging.html.

ble (Scholl et al. 2001, 170–72, fig.3). So subjects watch the now merged boxes move randomly around the screen, and in some merging conditions are able to select the boxes (the ends of the merged pairs) that were initially cued as targets. In other merging conditions (e.g., the rubber band condition) subjects cannot reliably select at the end which boxes (which ends of the merged pairs) were initially cued. As noted above, the inability to track in the rubber band merging condition suggests that the merged target and distractor boxes are no longer being grouped separately as distinct objects. Instead, the visual system treats the merged target-distractor box pair as a single object: it groups together all the features of the target box, distractor box, and rubber band around them as a single object. Call these merged target-distractor pairs *TD pairs*.

So, the target box ends in the rubber band merging condition provide an example of things for which no visual representation is constructed (see figure 4.3, which displays the TD pairs and their ends). Specifically, no visual representation is constructed while you do the MOT task involving the TD pairs. Since feature grouping is task dependent, the failure to track the target box ends only shows that visual representations of those target ends aren't constructed during the MOT task.

4.2.3 Visual Grouping Principles

The claim that the target ends of TD pairs are not represented by the visual system is supported by the failure of subjects to track these target ends in MOT tasks. But, you might object, tracking failure in MOT tasks need not mean that no visual representations of the targets are constructed. Perhaps representations are constructed, but the relevant grouping mechanisms which produce them don't enable multiple object tracking. For example, perhaps Pylyshyn is correct that there's a pre-attentive FINST-based grouping mechanism and the operation of this mechanism is required for tracking in MOT tasks. Then although tracking failure shows that this FINST-based mechanism doesn't group the visible features of the target ends of TD pairs (thereby constructing a visual representation), it leaves open that other mechanisms do group these features.

This objection can be pressed by suggesting two ways these other grouping mechanisms might come into play.[11] First, perhaps the visual representations of the whole TD pairs are built up out of visual representations of smaller component parts. On this suggestion earlier grouping mechanisms segment the scene

---

[11]Note that there is evidence against the detailed visual representations that would result from either process (see Poljac et al. 2012).
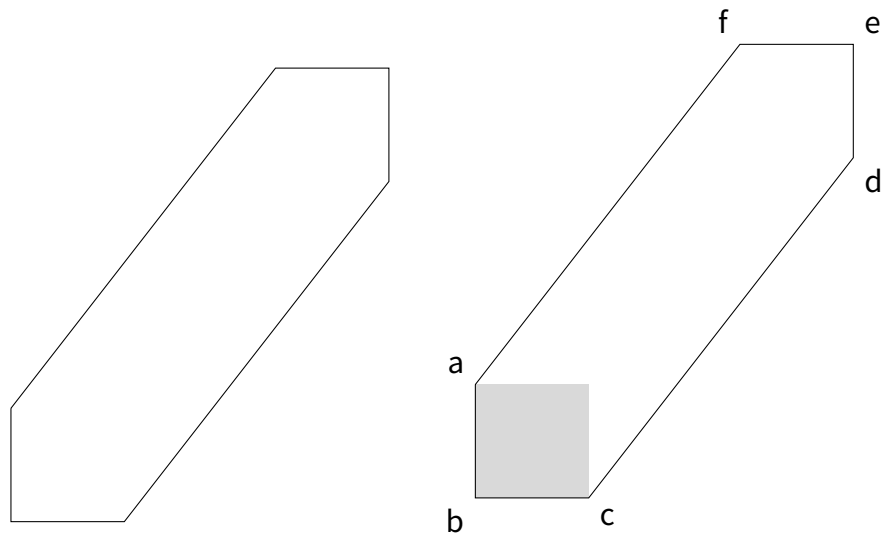
Figure 4.3: A single TD pair (left). The figure on the right labels the nodes of the TD pair and shades just the surface of the original target box. In constructing a representation of a TD pair the visual system groups together line segments ab, bc, cd, de, ef, and fa, along with the surface bounded between them. The target end of the TD pair to the left is the part corresponding to the shaded region and its adjacent line segments (lines ab and bc). The claim that the target ends are not represented in the visual system is the claim that these elements in the left TD pair (lines ab, bc, and the surface corresponding to shaded region on the right) are not grouped together when you look at it.

into more primitive object parts (including the target box ends) while a later, FINST-based mechanism groups these object parts forming a complex visual representation of the kind which can be tracked in MOT tasks. Second, perhaps the grouping mechanism which enables multiple object tracking operates early, segmenting the scene into large trackable chunks. Later, other grouping mechanisms decompose these chunks, forming more fine-grained representations of object parts (perhaps of the sort useful for other purposes, such as object recognition) (Hoffman and Richards 1984). That visual representations might have this complex nested structure is a common idea (Kahneman et al. 1992, 178; Scholl 2001, 18; Wagemans et al. 2012, 1188). For example, Marr's work on shape extraction posits that representations of shape are built out of representations of simple component shapes.

In a survey article Scholl (2001, 18) gives two lines of evidence for this complex, nested structure of visual representations. First, positing nested represen-

tations of component parts within other visual representations is theoretically useful. Marr ([1980](#), 211), for example, suggests that the nested structure of shape representations explains your capacity to identify objects (e.g., identifying a certain object as a hand, [Hoffman and Richards 1984](#), 67). This is because although the shape of some kinds of objects can vary greatly (e.g., a hand's shape depends on its articulation), the simple component shapes making them up remain constant. Second, results from Posner-style cuing experiments (another paradigm which, like MOT tasks, provides behaviorial, nonintrospective tests for grouping, see fn [9](#)) provide evidence that features from object parts are often handled as distinct groupings (results reported at [Singh and Scholl 2000](#), reference from [Scholl 2001](#), 18; see [Singh and Hoffman 2001](#) for a survey of part representation).

Do the visual representations of TD pairs constructed during the MOT task contain representations of their target ends? While there is no direct empirical evidence one way or the other, work on grouping principles suggest that visual representations of the target ends aren't constructed. *Grouping principles* describe the conditions under which features are grouped together. Presumably different sets of principles are associated with different grouping mechanisms, these sets describing their mechanism's operation. The study of these principles goes back to the gestalt psychologists, starting with Max Wertheimer's classic ([1923](#)) paper ([Wagemans et al. 2012](#), 1180). Wertheimer's classic grouping principles are listed in table [4.1](#). As Wagemans et al. outline in their survey article on gestalt grouping ([2012](#), 1181–82), since Wertheimer a number of other grouping principles have been discovered. These new grouping principles are listed in table [4.2](#), along with a few other grouping conditions reported by Scholl ([2001](#)) (see also [Spelke 1990](#)). Although Wertheimer originally supported these principles with simple demonstrations (see figure [4.4](#)), experimental techniques have been developed to more rigorously test grouping principles (see [Wagemans et al. 2012](#), 1182–88).

Given known grouping principles, it's unlikely that the target ends of TD pairs are represented. Representation of the TD pair target ends would require grouping the two residual line segments from the merged target box (lines ab and bc in figure [4.3](#)) along with the surface of just the target box (the shaded area in figure [4.3](#)). Such a grouping isn't supported by continuity and closure. These suggest that lines ab and bc should be grouped together with the other lines segments of the TD pair, lines cd, de, ef, and fa. Further, as the actual TD pairs are moving in the relevant case, and all these line segments are moving together, common fate and synchrony also suggest that all line segments are grouped together and not

1. ***Wertheimer's Classic Gestalt Grouping Principles:***

   Proximity: The closer two elements are to each other relative to other elements in the scene, the greater the tendency for them to be grouped together.

   Similarity: Similar elements (e.g., in color, size, and orientation) tend to be grouped together.

   Common Fate: Elements that move in the same way tend to be grouped together.

   Symmetry: Lines that are symmetric across an axis tend to be grouped together.

   Parallelism: Parallel lines tend to be grouped together.

   Continuity: Line segments tend to be grouped into continuous contours.

   Closure: Line segments tend to grouped into closed contours.

Table 4.1: List of Grouping Principles. Classic gestalt grouping principles taken from Wagemans et al. (2012, 1180–82). Note that some principles (e.g., continuity and closure) deal specifically with the grouping of certain types of item (e.g., line segments). Following Wagemans et al., for those principles which are general I refer to the grouped items as *elements* instead of as *features*. Also note that all principles have *ceteris paribus* clauses ("tend to be") because the extent to which features/elements are grouped according to them depends on what other grouping principles might be in play and global details of the stimulus.

just the two residual segments. Common region, element connectedness, and uniform connectedness do not support the target box surface being handled as its own element (grouped with the residual target box line segments). Instead they suggest the target box surface will be grouped with the rest of the surface within the whole TD pair. The upshot is that no known grouping principles provide support for the idea that the target ends of TD pairs are represented in the visual system.[12]

4.3   Step 2: Voluntary Attention

Using introspection, in this section I argue that the target ends of TD pairs are available for voluntarily attention. This is the first step in arguing that the tar-

---

[12]Even if some of the grouping principles did suggest that representations for target ends of TD pairs are constructed, there would still be the question of whether the grouping mechanisms described by those principles are operating during a MOT task in which you track the target ends of TD pairs. There would also be the question of what strength they had relative to the FINST-based mechanism that facilitates tracking.
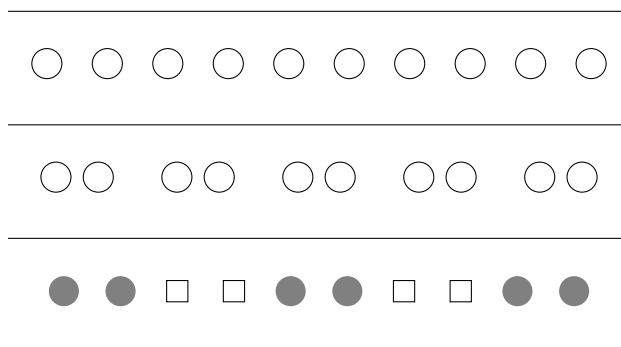
Figure 4.4: Demonstration of Select Grouping Principles. The top row should appear to be to a single unsegmented row of circles. The second row demonstrates proximity and should appear to divide into five pairs of circles. The last row demonstratives similarity and should appear to divide into three pairs of solid gray circles and two pairs of boxes. Adaptation of a common demonstration first given by Wertheimer (1923). Demonstrations reproduced by Wagemans et al. (2012, 1180).

get ends of TD pairs are consciously seen, i.e. that you have visual experience of them. Before beginning, note that introspection of your visual experience suggests that you have visual experience of the target ends. For example, when I look at the TD pairs (figure 4.2b), it seems to me that I don't just see the whole TD pairs, but also their target ends. But you might deny these introspection-based intuitions: perhaps it seems to you that you have a visual experience of the whole TD pairs, but not of the target ends. After all, the objection continues, you often have a visual experience of a whole object without having visual experiences of all its parts (e.g., occluded parts and the backside). The argument I develop in this section and the next (§4.4) avoids this objection by avoiding direct appeals to introspection of visual experience.

Before proceeding, some clarifications on voluntary attention are required. Attending is a person-level action directed at perceived things.[13] Three examples include visually tracking an object as it moves past you, visually searching a scene for a specified object, and being startled by a bright flash. In the first case you hold attention on the moving object, in the second you shift attention around the scene, and in the last case attention is grabbed by the flash. Just as the person-level action of seeing (or having visual experiences) can be distinguished from the underlying subpersonal, functionally-defined subtasks which enable it, attending should be distinguished from its underlying subpersonal, functionally-

---

[13]This description of attention given in this paragraph is taken from another of my papers.

2. *New Gestalt Grouping Principles:*

Generalized Common Fate: Elements that change their luminance value together tend to be grouped together (Sekuler and Bennett 2001).

Synchrony: Elements that change together, along any feature, tend to be grouped together, even if the changes aren't in the same direction (Alais et al. 1998).

Common Region: Elements that lie within the same bounded area tend to be grouped together (Palmer 1992; Palmer and Beck 2007).

Element Connectedness: Distinct elements that share a common border tend to be grouped together (Palmer and Rock 1994; Palmer and Beck 2007).

Uniform Connectedness: The visual system initially segments the scene into mutually exclusive connected regions with uniform, or smoothly changing, features (Palmer and Rock 1994).

3. *Misc Grouping Conditions:*

Minimize Coincidences: Elements will be grouped in a way that minimizes coincidences. E.g., two colinear, coterminous lines will be grouped a single contour instead of being treated as distinct lines that happen to align (Feldman 1999).

Occlusion: Represented objects (grouped sets of elements) survive static and dynamic occlusions (Behrmann et al. 1998; Moore et al. 1998; Scholl and Pylyshyn 1999).

Contour Closure: The closure the resulting contour is neither necessary nor sufficient for a set of line segments to be grouped (Avrahami 1999).

Uniform Connectivity: The grouping of two elements requires some uniform or continuous connection between them (Watson and Kramer 1999).

Table 4.2: List of Grouping Principles (Continued). New gestalt grouping principles taken from Wagemans et al. (2012, 1180–82). Misc. grouping conditions taken from and Scholl (2001, 30–32). Following Wagemans et al., I refer to the grouped items as *elements* instead of as *features*. Also note that all principles have *ceteris paribus* clauses ("tend to be") because the extent to which features/elements are grouped according to them depends on what other grouping principles might be in play and global details of the stimulus. Citations following a principle refer to studies that support the principle and are taken from the Wagemans and Scholl articles.

defined subtasks and the neural processes which realize them (Mole et al. 2011b, xi; Allport 2011, 26; Watzl 2011, 147; Wu 2011a, 106, 2011b, 97). Call these underlying neural processes *attentional systems*. Crucially, attending can be voluntary (endogenous) or involuntary (exogenous). The examples given above are the

starting point: in the first two it seems that the attention is under your control, or being deliberately guided. In the last case—of sudden bright flashes—the way attention is grabbed is not under your control. In these cases attention looks more like an automatic reflex. Substantial empirical evidence suggests that different attentional systems underlie the two forms of (personal-level) attention (Carrasco 2011, 1488–9).[14]

4.3.1 The Introspection Result

I claim that when you look at an object, it seems to you that you can voluntarily attend to any part of it which you can identify as a part. To support this claim about the possibility of attention to identified parts, consider again the TD pairs. When I look at the TD pairs in figure 4.2 (or watch the video of them moving) it seems plausible that I'm able to pick one TD pair and voluntary attend to its target end. At least, I can do this once the target end is identified to me as a potential target (e.g., as it is on the right in figure 4.3). This introspective claim can be further supported by noting that it's very easy to track the target end of a single TD pair, a task which seems to require voluntarily attending to that target end. Note that the failure to successfully track all four target ends in this MOT task does not show that you cannot voluntary attend to the target ends. At best, failure to track all four target ends in this task only shows that you cannot voluntarily keep attention directed to all four at once—it doesn't show that a single target end cannot be voluntarily attended.

Consider a second example, the odd shaped object in figure 4.5a. Figure 4.5b highlights a part of this object that's consciously seen as a part, i.e. for which a visual representation is plausibly constructed. Figure 4.5c explicitly renders an arbitrary part that, when looking at 4.5a, isn't plausibly seen as a part. (Figure 4.5d highlights this arbitrary part within the object.) It's easy to voluntarily attend to the part of 4.5a that's highlighted in 4.5b. Crucially, I claim, once the arbitrary part in 4.5c/d is identified, it's almost as easy to look at 4.5a and voluntarily attend to that part. When I look at 4.5a it introspectively seems to me that I'm able to voluntarily attend to the part from 4.5c/d.

---

[14]It should be clear that I assume that attention to objects and their parts is possible. I will argue for the further claim that voluntary attention to the target ends of TD pairs is possible. Substantial empirical evidence supports the starting assumption about the possibility of attention to objects and their parts (Scholl 2001, 5–13).
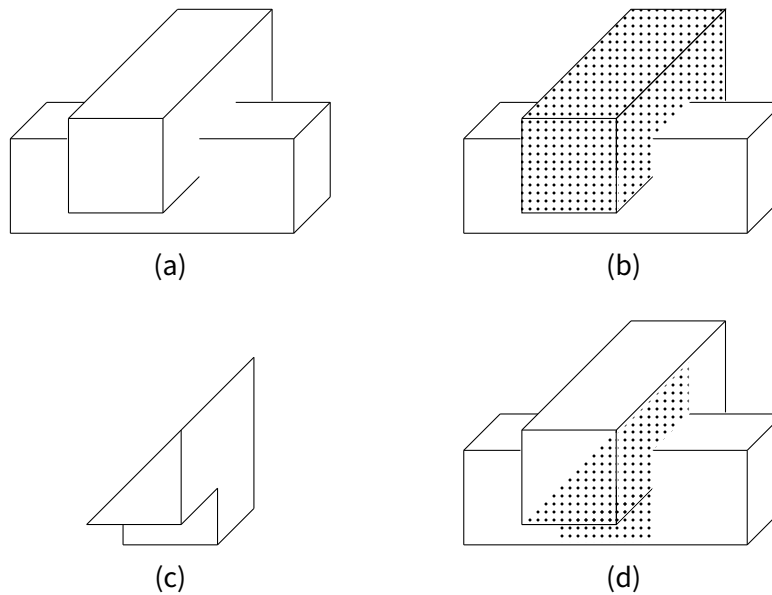
Figure 4.5: Contrast the part of the object (a) that's highlighted in (b) with the part presented in (c) and highlighted in (d). The part highlighted in (b) is seen as a distinct part, suggesting that the visual system automatically groups its features into a part. In contrast, the part in (c) is not automatically seen as a distinct part. But once it's presented, this part can be recognized in the object (a).

4.3.2 Voluntary Attention and Identification

But how could you voluntarily attend to something not represented by the visual system? At any time there will be some number of objects and their parts in within your field of view. The features of some of these will be grouped together, or represented, by the visual system: you will consciously see them as distinct objects and object parts. Although a deeper explanation in terms of attentional systems is required, presumably these objects and objects parts are immediately available for attention. But attending to an object or object part the features of which are not grouped together by the visual system requires first *identifying* the object or part. Presumably this is a conceptually mediated capacity which requires having the concept OBJECT or OBJECT PART, or at least some kind of pattern or representation of objects against which you can compare the set of features. The concept or pattern that is deployed in specific instances of identifying is usually more fine grain, e.g. BOX or TARGET END. Crucially, once you identify a set of visible features as features of an object or object part, that object or object part can be attended just the same as if the visual system had grouped its

features together. While this is an empirical claim, introspection provides some preliminary defeasible evidence for it. Once the contrast is seen between things that the visual system has automatically picked out and things that are identified, it introspectively seems almost as easy to attend to the identified ones as it is to attend to the ones for which the visual system has constructed representations.

It makes sense that identified objects and object parts can be voluntarily attended. When voluntary, attention is guided by nonperceptual cognitive states like beliefs, desires, goals, and—most important—intentions. On one view, intentions are propositional attitudes the contents of which are "plans" that represent actions (Mele 2009, 692). The intentions that guide voluntary attention have contents representing the subject attending to some visually experienced thing. Presumably the representation of what's visually experienced by the visual system suffices to form intentions with this content. Similarly, it's not implausible that post-perceptual conceptual identification of what's visually experienced also allows for intentions with contents that represent the subject attending to the identified thing.

To summarize, what you can voluntarily attend is not constrained by how your visual system segments a scene, i.e. by what is represented by the visual system. A high-level, conceptually mediated capacity for identification plus the intention-guided character of voluntary attention allows you to voluntarily attend to objects for which the visual system does not construct representations. Note that this explanation of how you can attend to objects and object parts for which no visual representations are constructed explains why tracking multiple target ends of TD pairs isn't possible, despite the possibility of tracking a single target end. If visual representations of the target ends were constructed, then attention to (and hence tracking of) them would would relatively unconstrained and effortless. But, as seems to be the case, no such representations are constructed and attending to the target ends requires post-perceptual conceptual processing. But this post-perceptual processing is relatively constrained (by, e.g., the limits of working memory, see Beck 2012; Hutchinson and Turk-Browne 2012). These constraints limit you to attending to only one of the target ends.

4.4 Step 3: From Attention to Visual Experience

The last two sections argued that the target ends of TD pairs are not represented by the visual system during MOT tasks, but are available for voluntary attention. This final section completes the overall argument by arguing that since the target ends of TD pairs are available for voluntary attention, you have visual experience

of them. The key claim is that visual experience of a thing is necessary for it to be available for voluntary attention. If you didn't have a visual experience of the target ends of TD pairs, then they would not be available for voluntary attention.

This claim about the necessity of visual experience for the availability of voluntary attention follows from a widely held view about the relationship between attention and perception. The view, often tacitly assumed, is that voluntary attention is a mechanism which operates to select consciously perceived things (Valberg 1992b, 21; Scholl 2001, 20; Pylyshyn 2007, 59; Levine 2010, 181; Dickie 2010, 216, 2011, 303 Wu 2011a, 109). On this view attention and conscious perception are different personal-level mental acts and their underlying subpersonal systems are functionally distinct. Objects and object parts are first perceptually experienced (consciously perceived), and then attention is a distinct action which operates on those objects and their parts. If this view is correct, then what's available for voluntary attention are the things you consciously perceive. So in the case of vision, what's available are things you consciously see. So the availability of the TD pair target ends requires that they are consciously seen.

Note that some alternative views on the relation between perceptual experience and voluntary attention integrate them in ways that block this last step. For example, Prinz (2011) says that to consciously see something is to have a visual representation of it selected by subpersonal attentional processes which make it available to working memory. Of course, this assumes the account of seeing which I reject in this paper, but we can abstract away from the details and say that Prinz holds that attending to something seen is both required, and sufficient, for you to have a visual experience of it. If this view is correct, then consciously seeing the TD pair target ends cannot be required for their availability for voluntary attention. If attention is required for experience, then unattended things—like the TD pair target ends in the MOT task—aren't experienced.

I won't defend the standard view over Prinz's more integrative view here. But you might object that there's empirical evidence that visual experience is not required for the availability of voluntary attention. Robert Kentridge (2008; see also 2011) has done Posner-style cuing studies with the blindsight patient GY, which, he argues, show that GY attends to things for which he has no visual experience. If GY can attend to things he unconsciously sees, then experience of them was not required for their availability for attention. In these studies arrows are flashed in the intact portion of GY's visual field which point to areas in the blind field. Then a vertical or horizontal line is flashed in the blindfield and, as in standard blindsight tests, GY must guess the orientation of the line. Accuracy in

the guessing task is better when location of the line is congruent with the direction of the arrow. The posited explanation is that the flashed arrow cues attention to the area in which it points, and this cued attention speeds reaction time. (Increased reaction time is a standard way to operationalize or measure attention in experiments.)

Kentridge's work does not pose a problem for the final step in the argument. First, there are alternative explanations of the results which don't posit that GY is attending to areas of his blindfield. Prinz discusses several of these (2011, 193–94). But even if Kentridge is correct that there is attention involved, there's only a problem on a very strong and implausible interpretation. Specifically, there's only a problem if GY is voluntarily attending to the lines themselves. If the attention is involuntary or not directed at the lines themselves (and instead, e.g., is directed at spatial locations), then there's no problem. Both of these are plausibly true. First, like any other cuing task, the kind of attention involved in Kentridge's studies is presumably involuntary. It's not that GY sees the arrow and deliberately directs attention in the indicated direction. Instead, the cued arrow grabs attention.[15] Second, attention is being cued to spatial locations, not the lines themselves. Kentridge himself characterize the results this way, saying that "we had demonstrated selective spatial attention in blindsight" (2011, 239).

4.5   Conclusion

In conclusion, while performing the rubber band condition merging MOT task from Scholl et al. you consciously see not only the whole TD pairs, but the target ends of the TD pairs too. This is established by appealing to an ability to voluntary attend to them and the requirement of visual experience for that ability. That subjects fail to do this MOT task shows that the grouping mechanism behind MOT is not constructing a representation of the target ends. More general considerations about known grouping principles suggest that no grouping mechanisms produce representations of the target ends. Thus, the target ends of TD pairs, perceived during the MOT task, are an example of something consciously perceived but not represented by the visual system. The example shows that the construction of a representation of a thing by the visual system is not constitutive of consciously seeing that thing.

---

[15]GY claims he can direct attention to areas of his blindfield (Kentridge 2011, 239). But the setup in which the results were obtained is a cuing task.

Chapter 5:
The Role of Experience in Demonstrative Thought

5.1   Introduction

Visually perceiving an object—seeing it—allows you to think of it directly, in a way that doesn't require conceptualizing it or thinking of it via a name or a description (Pylyshyn 2001, 154, 2007, 14–8; Campbell 2002, 7; Martin 2002a, 178–81; Raftopoulos and Müller 2006, 252,264; Dickie 2010, 213–15, 2011, 297; Jeshion 2010c, 134; Levine 2010, 185; Recanati 2010, 2012, 29).[1] Typical examples of thinking include mental acts such as judging, supposing, wondering, inferring, and occurrent instances of believing. For example, say you have a clear view of a red ball. Just by focusing your attention on it you seem to thereby be able to think thoughts about it. These thoughts would naturally be expressed with a demonstrative term like 'that', e.g. you might judge "*That*'s red", wonder "How much is *that*?", or believe "I'd like to purchase *that*" (all said while pointing to the ball). These are often called *demonstrative thoughts* (e.g. Smithies 2011b, 7), but I will call them *vision-based thoughts* to avoid suggesting that the thoughts themselves share controversial features with the demonstrative speech acts that express them.[2]

Seeing is typically conscious: it typically involves a visual experience of the seen object. Often this is referred to as 'phenomenal consciousness'; you (in alternative jargon) are *phenomenally conscious of* the seen object, you *consciously* perceive it visually, or are *visually aware* of it. You have a visual experience of (for

---

[1]Specifically, seeing an object normally allows you to *select* it for thought. For example, in simple judgments you both select an object and attribute a property to it (see Levine 2010, 178). Although thinking of an object involves conceptualizing it in the sense that you attribute some property to it, seeing allows you to select without deploying a concept.

[2]For example, you might think that vision-based thoughts involve demonstrative-like mental representations, similar to how the demonstrative speech acts use demonstrative terms (e.g., Bach 2010, 55; Levine 2010, 179; Recanati 2012, 57–67). But this is controversial and need not be assumed (see Millikan 2012).

example) a red ball when (1) you see it and (2) there's "something it's like" to see it (Nagel 1974), or (alternative put) there's some way the red ball is (or *looks*) to you visually. Visual experience involves a subjective "felt quality". To help further narrow down the experiential aspects, we can note that your visual experience of the ball (what it's like to see it) changes depending on the lighting, your relative position, and other contextual factors. Likewise, consciously seeing two objects that differ in color or shape (to give two examples) will, unsurprisingly, lead to visual experiences that differ in corresponding ways. It can also help to contrast visual experience with two other varieties of consciousness. Seeing might be conscious in the sense that you have access to its content (access consciousness, see §5.4.1) or in the sense that you are aware of, or know, that you're seeing (self or higher-order consciousness). Visual experience is not merely seeing that's access or self conscious.

A plausible suggestion is that visual experience plays some functional role in vision-based thought. This is plausible, at least on a first pass, because vision-based thought does not seem possible in known cases of unconscious seeing, specifically blindsight (see §5.2). In this paper I give a proposal for this functional role. The proposal has two parts:

Part 1: In many cases voluntary focal attention plays the target-setting role in vision-based thought: you often set the targets of vision-based thought by voluntarily and focally attending to them.

Part 2: The functional role of visual experience is to make visual information available for use in the voluntary control of focal attention. Call this the *attention-guiding role*.

It's both prima facie plausible and widely accepted that the targets of vision-based thoughts are often set by voluntarily and focally attending to them (e.g., Campbell 1997, 2002, 2004; Levine 2010; Wu 2011a,b; see also Raftopoulos and Müller 2006, 253; Dickie 2010, 234; Jeshion 2010a, 1). After developing this part of the proposal (§5.3.1) I take it for granted. My main aim is to present both introspective (§5.3.2) and neurophysiological (§5.3.3) evidence that visual experiences play the attention-guiding role.[3]

---

[3]The more general idea that experience has the functional role of making perceptual information available for use in cognitive functions (e.g., in reasoning and action control) is not new. For example, the idea is found in Kriegel (2004, 184–85) and Milner and Goodale (2006, 222; although see Brogaard 2011b). Note that this is different from the idea that experience arises from, or is explained by, availability for use in cognitive functions (e.g., Baars 1988; Tye 1995, 137–43; Prinz 2011, 184–87).

A related question is whether a super blindsighter—a person who, in some patch of her visual field, sees objects without having a visual experience of them but still is access conscious of visual information from that patch—can have vision-based thoughts about objects in her blindspot. The question is whether, given the functional role of visual experience, vision-based thought is possible when you see an object without any accompanying experience of it, but yet are still access conscious of the relevant visual information. Often the question is put less precisely by asking whether visual experience, or phenomenal consciousness, is *necessary* for vision-based thought. Intuitions about what you would be capable of if you were a super blindsighter suggest a negative answer: vision-based thought is not possible in such a case. I argue (§5.5) that on the proposal that experience plays the attention-guiding role there are good reasons to deny these intuitions. Vision-based thought is possible with only access consciousness, so long as the mode of accessing the visual information allows for its use in the voluntary control of focal attention. This is in contrast to other proposals for the functional role of experience (specifically Campbell 2002, 7 and Smithies 2011a, 264, 2011b, 7,19), which (at least it's argued) make it so that visual experience is necessary for vision-based thought.[4]

The structure of this paper is as follows. This section concludes with three important points about the project itself. Then §5.2 explains nonconscious seeing. Next, §5.3 develops the two-part proposal for the role of visual experience in vision-based thought and presents the introspective and neurophysiological evidence that visual experience plays the attention-guiding role (part 2). A relatively direct empirical test is also suggested. After that §5.4 replies to two objections to part 2 of the proposal. The first is that there are purported cases of experience (i.e., phenomenal consciousness) without access consciousness. This is a problem since part 2 is, essentially, that visual experience has the role providing access consciousness to visual information for the purpose of voluntarily guiding attention. The second objection is aimed at my use of neurophysiological evidence and is that there's a two-step gap between experience-related neural activity, the visual experiences themselves, and their phenomenal properties (see Kriegel 2004, 174–75). The objection suggests that the neurophysiological evidence from §5.3.3 only supports that neural activity, but not that expe-

---

[4]Although they don't develop proposals about the role of visual experience in vision-based thought, Valberg (1992b, 21), Johnston (2006, 263–65), Siegel (2006a), and Dickie (2011, 294,298) suggest that visual experience is necessary for it. In contrast, Kelly (2004, 283–4) and Wu (2011b, 115–18) rejects this necessity claim.

riences or their phenomenal properties, has the attention-guiding role. Finally, §5.5 turns to the question of whether visual experience is necessary for vision-based thought. It combines the two-part proposal from §5.3 with empirical work to argue that it is not, at least in the sense that access consciousness may suffice for vision-based thought.

The first important point is about the evidence given below (§5.3) for the functional role of visual experience in vision-based thought. It is meant to be merely supportive or suggestive. The idea is not that the proposal for visual experience's functional role logically follows from, or is entailed by, facts about introspection or the neurophysiological structure of the visual system. The second point is that although the discussion is framed in terms of seeing *objects*, it's meant to apply more broadly to seeing properties, events, locations, and whatever other kinds of things are perceptible through vision. The third point is that although this discussion is about vision, it's hoped that most of the points here generalize to other sensory modalities. For example, I think it's plausible that voluntary focal attention sets the targets for any kind of perception-based thought and that, in all sensory modalities, perceptual experience fills the role of providing perceptual information for use in the voluntary control of attention.

5.2 Preliminaries: Actual Blindsighters and Nonconscious Vision

Before discussing the two-part proposal for visual experience's role, it's worth saying more about nonconscious seeing. This, presumably, is just seeing without accompanying visual experience of the seen object. But it might be difficult to imagine what this would be like, or if the idea is even coherent. (If there's no visual experience, then in what sense is the object seen?) It turns out there are individuals with a condition, called blindsight, giving them something very much like nonconscious vision (see Brogaard 2011a; Brogaard 2012 for careful discussion).[5] Blindsight gives both a concrete example of nonconscious seeing and suggests a general way to conceptualize it. As mentioned above, the intuition that blindsights lack the ability to have vision-based thoughts provides initial motivation for thinking experience has a role at all.[6]

---

[5]More controversial examples of nonconscious seeing include unilteral neglect, attentional blink, and inattentional blindness (Prinz 2011, 177–78).

[6]Blinsighters lack both experience and access consciousness (see §5.4.1). This makes placing real weight on considerations about blindsighters problematic. The lack of access consciousness might prevent vision-based thought. Still, the case provides some *initial* motivation.

What is blindsight? In humans and other primates about 90% of the ganglion cells projecting from the retina synapse into the dorsal lateral geniculate nucleus, which in turn projects into V1, the primary visual cortex (Weiskrantz 2009, 69). This is called the central visual or geniculo-striate pathway. V1 itself is just the start of a large network of interconnected cortical areas responsible for processing visual information. In humans the destruction of some portion of V1 leads to total blindness (the absence of visual experience) in the part of the visual field to which that portion of V1 maps topographically. But in cases in which the damage is restricted to V1 (leaving the areas of visual processing beyond that intact) some residual visual functioning often still remains, leading to a loose collection of abilities. *Blindsight* cases are ones that involve this residual visual functioning despite a lose of visual experience from V1 damage.

For example, in some experiments D.B. (one of the original people in whom blindsight abilities were discovered) was shown a flash in the blind patch of his visual field and was asked to either direct his gaze or point to the location he *guessed* the flash had been. This he could do well above chance (Weiskrantz 2009, 87–93). Similarly, in another set of experiments D.B. was found to be able to guess accurately the orientation of lines (vertical or horizontal, vertical or slanted, an "X" or an "O") that were flashed in his blindspot. These and similar results have been replicated many times both with D.B. and with other patients with blindness caused by V1 damage. When prompted to guess from a limited set of choices these blindsighters can reliably identify the orientation, location, spatial frequency, wavelength, movement, flicker and (at least in D.B.) the form or shape of objects in their blindspots (Weiskrantz 2009, 17–21,53).

The standard interpretation of blindsight is that the residual visual functioning is explained by the small portion of the optic nerve that lies outside the central visual pathway. Ganglion cells from the retina outside this pathway project to extrastriate cortex (V2–V5), inferior temporal (IT) cortex, and subcortical regions via 9 or 10 other pathways (Weiskrantz 2009, 59,69). These pathways bypass the damage in V1 and instead project directly to "higher" areas of visual processing. The idea is that V1 must be necessary for conscious seeing (seeing with experience of the seen object), but these additional pathways from the retina to higher visual areas facilitate or allow for some level of visual processing that, while not leading to conscious seeing, still allows for some limited forms of functioning. More abstractly, it seems that while objects within the blindspot aren't consciously seen, there is still visual states which represent those objects and have some influence over behavior (e.g. Block 1995, 230). This provides a general way

to understand nonconscious seeing: it happens when there's some level of visual processing of sensory input (leading to what count as visual states with representational content), but that processing for whatever reason does not lead to visual experience of the seen objects.

Now, imagining what it would be like to be a blindsighter, insofar as that's possible, suggests that it's not possible to have vision-based thoughts about objects in the blindspot (Campbell 2002, 7; Johnston 2006, 264–265; Smithies 2011b, 6,26; cf. Kelly 2004, 283). Imagine, for example, that there is an object in front of you that—as it would be natural to say—you could not *see*. More precisely, you have no visual experience of it. In this case it seems clear that you cannot have vision-based thought about the object. Now imagine that, like a blindsighter, given choices about basic features of the object and asked to guess you reliably guess the correct answer. Perhaps this happens because, like in blindsight, some amount of visual processing is restored. It's plausible that gaining this ability does not make it so that you can now have vision-based thought about the object.

5.3   The Role of Visual Experience

5.3.1   Part 1: Voluntary Attention

Part 1 of the proposal is that in many cases of vision-based thought voluntary focal attention plays the target-setting role. That is, you often set the targets of vision-based thought by voluntarily and focally attending to them. At any one time you usually are seeing more than one object. But when you exploit vision to think of a seen object in a vision-based thought you manage to single out one of the multiple seen objects as the target of thought. Part 1 of the proposal says how you achieve this: you single out or think of a single seen object by attending to it. *Focal* attention in particular is specified as the target-setting mechanism because attending to multiple objects at once (as in *divided* attention) would not manage to single out an object as the target.[7] This much I take to be prima facie plausible (if not likely) and widely accepted (e.g., Campbell 1997, 2002, 2004; Levine 2010; Wu 2011a,b; see also Raftopoulos and Müller 2006, 253; Dickie 2010, 234, 2011, 294; Jeshion 2010a, 1; although see Levine 2010, 178 for a critical discussion).

Focal attention to a seen object is either voluntary or involuntary. Part 2

---

[7]From a broader point of view part 1 of the proposal is putting forward a metasemantic view on the content of vision-based thought. The basic metasemantic question about vision-based thought is: when you have a vision-based thought about some seen object O, what makes it the case that O is the target of that vision-based thought? The answer, according to part 1 of the proposal, is just that it's in virtue of your voluntary focal attention to O that it is the target.

of the proposal specifies that visual experience plays the attention-guiding role when the focal attending which sets the target of vision-based thought is under voluntary control. So it's important to get clear on the distinction between voluntary and involuntary attention. Attention, whether focal or divided, can be allocated to objects (and properties of objects and locations) in the visual field in one of two ways: you can deliberately direct attention to a seen object, or an object might involuntarily "grab" attention. An object might grab attention because of high salience or sudden movement. The distinction can be brought out ostensively through examples. The first case, usually called *endogenous* or *voluntary* attention, is exemplified in visual search tasks. For example, you might voluntarily shift attention between objects (e.g., a mix of red and blue 'A' and 'B') in a display looking for those with some distinguishing features (e.g., blue 'B'). The second case, usually called *exogenous* or *involuntary* attention, is exemplified in cases in which you are distracted by a sudden bright flash or loud bang. Introspection suggests that you can voluntarily *shift* attention between objects in your visual field, *hold* attention on an object (e.g., while tracking it), *divide* attention between objects, and *zoom* attention between larger and smaller sections of the visual field.

Before giving evidence for part 2 of the proposal, note that part 2 leaves visual experience no role to play when the target of vision-based thought is set through involuntary focal attention. Elsewhere I've argued that the targets of vision-based thought can only be set through voluntary attention. The main problem with involuntary attention is that it's often divided between multiple seen objects (or their properties and locations) even when vision-based thought intuitively seems possible. But even if the targets of vision-based thought can be set through involuntary attention there's no problem for part 2 of the proposal. It would just turn out that visual experience only plays a role in certain cases of vision-based thought: those in which the focal attending which set the target was voluntary.

5.3.2 Part 2: Introspective Evidence

In this section and the next (§5.3.3) I will present two lines of evidence that visual experience fills the attention-guiding role (part 2 of the proposal) along with a suggestion (§5.3.4) for how this proposal might be directly tested in future empirical work. The first is introspective evidence. Introspecting what it's like in some clear cases of voluntary attention strongly suggests that you rely on information available from our visual experiences. Consider two examples. The first

is visual search of the kind just mentioned. In searching through a display of (for example) a mix of red and blue 'A' and 'B' your attention shifts from one item to another. But each time you make a shift, it seems, you direct the shift to the next item using your visual experience. What it's like to make the shift is that you have a visual experience of (say) a red 'A' being a few degrees in some direction from your current focus of attention and shift to *that*, to the experienced red 'A'. You know the direction in which to shift, when to stop, and how to hold attention all based on your visual experience of the 'A'. The second example is voluntarily tracking an object, say a red ball as it rolls past you. Tracking the rolling red ball requires holding attention on it. But it seems, introspectively, as if you rely on your visual experience of the ball to hold attention on it.

This introspective evidence is something like an existence proof. The idea is to present cases in which—as it turns out—you use information from visual experience to voluntarily guide attention. Introspecting what it's like to voluntarily guide attention in the above two examples reveals your use of experience-based information. In them you introspectively find yourself using experience to guide attention. But what if you engage in visual search and visual tracking tasks, of the kind just described, but don't introspectively find any use of experience? In that case Susanna Siegel's method (2007) of using contrast cases provides an indirect way to use introspection to support the attention-guiding role of visual experience. As applied here,[8] in this method a pair of cases are presented which everyone can agree differ introspectively in a specific way. Then it's argued that the best explanation of that agreed introspectable difference presupposes that (in at least one of the two cases) experience provides information you use to guide attention. If successful, then the contrasting cases along with the best explanation of their introspectable difference support the use of experience-based information in guiding attention in at least one of the cases.

The contrasting examples I will use each involve a simple attentional shift between dots in figure 5.1. For these tasks the array should be at a distance which puts the inner circle with its dots only a few degrees of visual arch from the center dot. So, the dots on the inner circle should be comfortably visible while the dots on the outer circle are towards the edge of your visual periphery. The contrasting

---

[8]Siegel uses her method of contrast cases to argue for claims about the content of visual experience. My use of it here involves some adaptation from its original form.
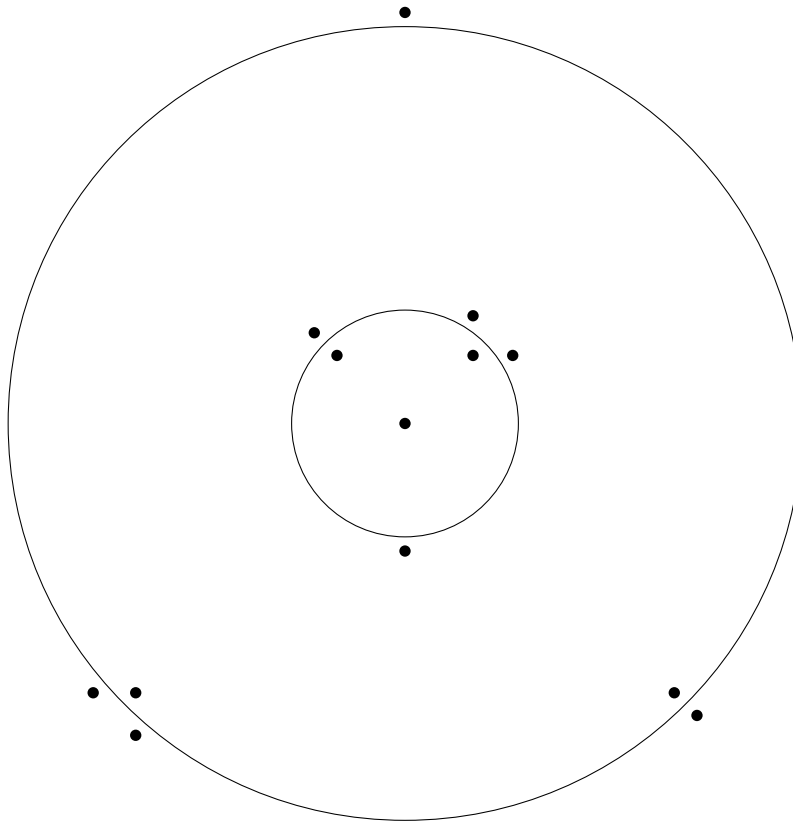
Figure 5.1: Array of dots for tasks 1 and 2.

examples or tasks are:

Task 1: Foveate the center of the array and hold attention on the center dot. Pick a dot on the inner circle and shift your attention, either covertly or overtly, between the center and that dot.[9]

Task 2: With the array still at the same distance from you, this time pick a dot on the outer circle and shift your attention, either covertly or overtly, between the center and that dot.

The introspectable difference between the two tasks—the way they contrast—is their difficulty. Everyone should agree that shifting attention from the center to a nearby dot (task 1) is easy, while shifting attention from the center to a distant

---

[9]Cover shifts of attention involve shifting attention without shifting gaze, i.e. without changing the area you're foveating.

dot on the periphery (task 2) is difficult.[10] More specifically, the resolution with which you can localize new targets of attention goes down as you go out to the periphery. For example, it's easier to shift attention to the top single peripheral dot than to shift attention to one of the three dots in the lower-left peripheral dot cluster.[11]

The only two plausible explanations for the introspectable difference between tasks 1 and 2 both presuppose that experience is used to guide the shift in task 1. The first is that in both tasks you rely on your visual experience of the dots to select the targets and shift attention. Visual experience itself of objects in the periphery is much less detailed and precise than visual experience at the center of your field of view. It's this degradation of visual experience, and associated loss of experience-based information at the periphery, which explains the increase in difficulty from task 1 to task 2.[12] The second is that tasks 1 and 2 involve using different sources of information to make the shift. Task 1, the short-range shift, relies on experience-based information. Task 2, the long-range shift, relies on information stored in memory to make the shift. Since stored information about the peripheral dots' locations is less precise than experience-based information, the long-range shift in task 2 is more difficult.[13]

---

[10]Note that it's easy to shift attention in the rough direction of the peripheral dots. With the peripheral dots then in center view you can easily attend to one of them. But this shift-then-select strategy isn't task 2; task 2 is to shift attention directly to the peripheral dots.

[11]Of course, both the general claim about the relative difficulty of tasks 1 and 2 and the specific claim about the resolution of localizing targets of attention are open to empirical testing. But introspection yields robust enough differences to serve the purpose here.

[12]Thanks to Casey O'Callaghan for suggesting this explanation.

[13]What are the alternative explanations? Any explanation of the difference in difficulty will need to identify the source of information used to make the shifts. There only seem to be three options: (1) information from visual experience, (2) nonconscious visual information (e.g., the kind processed in the dorsal stream and used to execute most visually guided movements), and (3) previous visual information stored in memory. If the same source of information is posited in making both shifts, then something about that source would need to explain why the long-range is more difficult. Experience-based information seems to be only one of the three that meets this constraint. If different sources are used, then explaining the difference in difficulty would seem to require that the information source used in the long-range shift (task 2) is less detailed than the one in the short-range shift (task 1). That leaves using nonconscious visual information for the short-range shift and experience-based information for the long range shift as the only alternative. But it just seems implausible.

5.3.3 Part 2: Neurophysiological Evidence

The second line of evidence comes from work on the neural correlates of visual experience and attentional control in Macaque monkeys (see Chalmers 2000 on neural correlates). The cortical regions in the brain primarily responsible for the voluntary control of attention are the frontal eye fields (FEF). There is a topographic mapping between the ventrolateral-to-dorsomedial direction of the FEF and the amplitude, or size, of attentional shifts. Roughly, the more dorsomedial regions of the FEF control larger amplitude attentional shifts (e.g., attentional shifts to the periphery of the visual field) while more ventrolateral regions control smaller amplitude attentional shifts. Crucially, while both dorsomedial and ventrolateral FEF receive projections from the dorsal visual stream, the ventral visual stream projects only to the ventrolateral regions. Because the ventral stream is associated with conscious seeing, its projecting into the FEF suggests that conscious seeing does provide visual information for use in the voluntary control of attention. The absence of projections from the ventral stream into the dorsomedial regions of the FEF further suggests that the second explanation offered above is correct. There is a difference in difficulty between tasks 1 and 2 because there's a difference in the sources of visual information which can be used to guide the shifts. Now for the details.

The frontal eye fields are defined as the areas of cortex in which low-current electrical stimulation ($\leq 50\mu A$) produces saccadic eye movements (e.g. Bruce et al. 1985, 714). In each hemisphere the FEF is located in Walker's areas 8A and 45 (Walker 1940), within the prearcuate sulcus (Stanton et al. 1989, 416) (see fig. 5.2). Bruce et al. (1985) suggest that its location is restricted to area 45 and only a portion of area 8A, what they call area 8Ac. The overall structure and activity of neurons in the FEF support the conclusion that it's one of the primary regions of cortex involved in the voluntary control of eye movements (saccades), but not involved in involuntary eye movement (for one modern study see Bruce and Goldberg 1985). The dorsomedial regions of FEF (area 8Ac) are associated with large-amplitude saccades, while the ventrolateral regions (area 45) are associated with small-amplitude saccades (Bruce et al. 1985, 714,730; Stanton et al. 1989, 426). There's no sharp divided between the two regions and their functioning; instead it's a continuous topographic mapping in which areas of the FEF more dorsomedial control longer amplitude saccades than those less dorsomedial.

More importantly, recent work strongly suggests that the FEF are involved

not only in the voluntary control of eye movements, but also in the voluntary control of attentional shifts (e.g., Schall 2004; Thompson and Bichot 2005; Thompson et al. 2005; Cohen et al. 2009a,b; Heitz et al. 2010; Lee et al. 2012; Ronconi et al. 2014; for reviews see Schall 2004, 1453–54; Armstrong 2011, 87–9). Some of the neurons in the FEF which respond before saccades seem to make a "salience map" of the visual field. Bruce and Goldberg (1985, 609) call these *visual cells*, since (1) they response before a saccade (to their receptive field) only if there's also a visual stimulus within their receptive field, and (2) they also respond to visual stimuli within their receptive fields even when there is no saccade.[14] As Thompson and Bichot argue (2005), unlike other areas of the visual system in which neurons respond selectively to a given feature (e.g., color or edges of certain orientations), these visual neurons within the FEF respond to visual stimuli based on their salience or behavioral relevance. The highest neural responses in these visual neurons corresponds to the most salient, or relevant stimuli. So the visual neurons within the FEF form a map of the visual field,[15] a given neuron responding to a stimulus within its part of the visual field (within its receptive field) when that stimulus is *attended*. The highest responses within this map are from neurons with receptive fields which cover the attended target. This salience map, along with other experimental results (cited above), makes the FEF a prime candidate as the main neural region responsible for facilitating the voluntary control of attention.

The primary visual cortex (V1) is the start of an interconnected set of visual processing areas which divide into two fairly distinct "streams": a dorsal stream running from V1 in the occipital lobe up into the parietal lobe and a ventral stream running down into the temporal lobe (Ungerleider and Mishkin 1982; Goodale and Milner 1992; Milner and Goodale 2006, 2008; Kravitz et al. 2011, 2013; see also Brogaard 2011b) (see fig. 5.2). The dorsal stream processes visuospatial information that's not directly part of the content of visual experience, but still used in the visual guidance of action. The ventral stream, on the other hand, processes object-form information that is part of the content of visual experience (see Milner and Goodale 2006, 221–28 for review). The highest regions of the

---

[14]In their important survey of 752 FEF neurons, Bruce and Goldberg (1985, 608) found that 409, or 54.3%, of the neurons had presaccadic responses. Of these presaccadic cells they studied 115 extensively, finding 46, or 40%, were specifically visual cells.

[15]As already noted, this map is topographically organized with respect to the ventrolateral-to-dorsomedial position of the visual neuron and the amplitude of the shift. But there's no global topographic mapping within the FEF for the direction of shifts. There is however some local organization, e.g. neurons close by each other tend to correspond with shifts of similar direction.
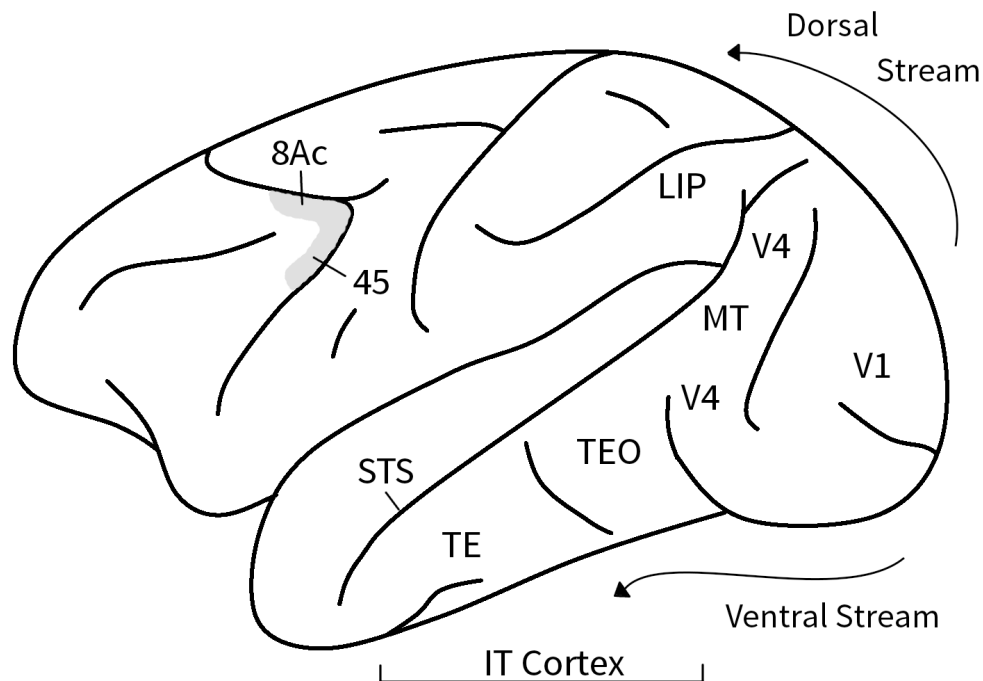
Figure 5.2: Diagram of Macaque monkey brain. FEF are shaded gray. Adapted from schematic brain outline made available by Ryosuke Niimi on http://rnpsychology.org/demo/index.htm. Labeling based on Schall et al. (1995, fig. 17) and Kravitz et al. (2013, fig. 1).

ventral stream in the IT cortex (areas TEO, TE, and the superior temporal sulcus, STS) are especially good candidate sites for neural correlates of visual experience. For example, in an important single-cell recording study of Macaque monkeys Sheinberg and Logothetis (1997, see Logothetis 1998 for review of related work) showed that the activity of 90% of recorded cells in these areas reliably predicted the visual experience of the animal in cases of binocular rivalry. In contrast, previous studies had showed that activity in only 18% of recorded cells in V1 and 20% and 25% in MT and V4 (respectively) predicted the animals visual experience (Sheinberg and Logothetis 1997, 3413; see also Crick and Koch 1995). Similarly, visual experience in binocular rivalry does not appear to be predicted by activity in the dorsal stream (see Milner and Goodale 2006, 225 for references).

The FEF receive and send projections to numerous cortical and subcortical areas. But what is most relevant here is the distribution of relatively direct pro-

jections into the FEF from cortical visual areas, since this gives some indication of the kind of visual information used in the voluntary control of attention. As noted above, the key result is that while areas from both the dorsal and ventral streams project to the FEF, the highest regions of the ventral stream (IT cortex) project only to the ventrolateral FEF (area 45) and not into the dorsomedial FEF (area 8Ac) (see Schall et al. 1995, 4466 and Kravitz et al. 2013, 41 for review). Areas from the dorsal stream, in contrast, project to both the ventrolateral and the dorsomedial FEF. For example, in one tracer study Shall and colleagues found projections into area 45 from ventral stream areas TEa, TE3, TEm, TF, and TEO (Schall et al. 1995, 4466,4484; see also Schall 1995, 70,76 for review).[16] The same study found that dorsal stream areas LIPv and LIPd project into both areas 8Ac and 45 (Schall et al. 1995, 4483, see also Schall 1995, 70 for review). No projections into area 8Ac from higher ventral stream areas (IT cortex) were found.

These results on the connections between the FEF and visual processing areas are important for two reasons. First, the projection of the highest regions of the ventral stream (IT cortex) into the FEF is significant. It means information from conscious visual processing is available to the visual neurons within the FEF which select targets of attention. So, it suggests that information from conscious seeing is available and used in the voluntary control of attention. Contingently, it could have turned out—but didn't—that only dorsal stream areas of visual processing projected into the FEF. Second, the difference between the projections into ventrolateral and dorsomedial FEF are congruent with the introspectable difference between tasks 1 and 2. At the neurological level there *is* some difference in the kinds of information available for short and long-range shifts. Interestingly, that the difference is a lack of projections from the ventral stream into the dorsomedial FEF supports the second explanation suggested above: information from conscious seeing is available for task 1 (the short-range shift), but not for task 2 (the long-range shift).

Note that the support of these neurophysiological considerations for visual experience having the attention-guiding role doesn't depend on strong assumptions about the neural correlates of visual experience. For example, although it must be assumed that the IT cortex (areas TE and TEO) is in some important way associated with conscious seeing, there's no need to assume that any given visual experience can be localized as activity of just a certain sort within some small region of IT cortex. Likewise, it need not be assumed that the content of

---

[16]These areas of IT cortex were also found to project to other sites outside the FEF, e.g. areas 8Ar and 12 (Schall et al. 1995, 4466,4484).

a given visual experience matches or is grounded in specific, localizable neural representations within the IT cortex (e.g., Prinz 2000, 249, 2006, 454, 2011, 174). The structure of connections between areas of visual processing and the FEF provides support for visual experience having the attention-guiding role even if the neural correlates of visual experiences are widely distributed and visual experience content can't be matched to localizable neural representations.[17]

5.3.4   Part 2: A Direct Test?

The proposal that visual experience plays the attention-guiding role might be more directly tested. If the proposal is on the right track, then we should expect inhibition of or damage to the relevant ventral stream areas of the IT cortex—the neural correlates of visual experience—to correlate with a loss of voluntary short-range attentional shifting. This inhibition might also be expected to leave the capacity for voluntary long-range attentional shifting intact. Although there is not, to my knowledge, any present studies along just these lines, there are studies of the effects of damage to V4 and (crucially) TEO, the posterior region of the IT cortex, on attention. For example, one study (Buffalo et al. 2005; see 141 for review of others) found that damage to V4 and TEO in Macaques lead to behavioral and neurological changes consistent with a loss in ability to focus attention voluntarily. Specifically, the damage nearly doubled object discrimination thresholds when there were distracters near the target but had no effect on object discrimination thresholds when there were no distracters. Similarly, the same study found that damage to V4 and TEO lead to a loss of attentional "filtering" effects in area TE (anterior IT cortex) typically observed during attention to an object surrounded by distracters.

5.4   Two Objections to the Attention-Guiding Role

The previous section provided both introspective and neurophysiological evidence for the proposal (part 2) that visual experience plays the attention-guiding role. This section responds to two potential objections: one related to Block's distinction between phenomenal and access consciousness, and one related to the distinction between neural correlates, experiences, and phenomenal properties of experiences.

---

[17]Thanks to Dan Burnston pointing out the need to make this point explicit.

5.4.1 Access Consciousness

An alternative way to put the claim that visual experience plays the attention-guiding role is by saying that visual experience provides what Ned Block has called access consciousness to visual states and their content. As Block defines it (1995, 231), a mental state such as the one you are in when you see an object is *access conscious* (or the state's content is access conscious) if, in virtue of being in that state, it's content is available for use (1) in reasoning, (2) in the rational control of action, and (3) in the rational control of speech.[18] Block's main point in distinguishing between experience (phenomenal consciousness) and access consciousness was to show that any proposal for putting experience in the functional role of providing access consciousness is a substantial thesis (Block 1995, 241–42). A loss of visual experience (as in blindsight) typically comes with a loss of access consciousness. Block wanted to cut off the further move, that many were prone to make, of saying that the lack of visual experience (i.e., phenomenal consciousness) explains the lack of access consciousness, i.e. that the experience had the role of providing thee access consciousness.

A preliminary point is that you might ask whether my proposal isn't making this same illegitimate move and inferring, from a lack of both phenomenal and access consciousness in blindsighters, that phenomenal consciousness (experience) has the functional role of providing access consciousness. The response to this objection is that it misconstrues the argument. I think it's plausible that blindsighters lack the use of visual information to voluntarily direct attention and that this provides some initial motivation for suspecting that experience has the role of providing access consciousness. But I'm not inferring the latter from the former. Instead, I'm relying—independent of any considerations about blindsighters—on the introspective evidence from contrast cases in §5.3.2 and the neurophysiological evidence in §5.3.3 to support the claim that experience provides access consciousness.

Moving on to the objection, Block has given cases that seem to involve visual experience without access consciousness. For example, Block suggests that in Sperling's now well known experiment (see Sperling 1960) participants are phenomenally conscious (i.e., have experience of), but not access conscious, of all the letters in the display jointly (Block 1995, 244, but see Prinz 2011, 190 for reply). Inattentional and change blindness (Rensink et al. 1997, 2000; Mack and

---

[18]Block says that (1–3) are jointly sufficient to make a mental state access conscious, but that (3) is not necessary.

Rock 1998; see Jensen et al. 2011 for review) provide other potential examples of experience without access consciousness. In these cases (the suggestion goes) you still have a visual experience of the unattended object or masked change, but fail to notice it because the visual state is not access conscious.

The objection presses that if these are cases of visual experience without access consciousness, then my proposal is wrong. The cases show that visual experience doesn't provide access consciousness. The reply is that, even assuming these are cases of visual experience without access consciousness, they don't show that visual experience fails to make visual information available for use in voluntarily control of focal attention. The crux is that access consciousness need not be all or nothing (Block 1995, 232). Consider inattentional and change blindness. It might be that visual experience of an unattended object affords a mode of access to the visual content that makes it available for the voluntary control of focal attention, but not available to the cognitive processes involved in these tasks, e.g. working memory. The basic idea is that a mode of access to visual content might make that content available for some tasks (e.g., the voluntary guidance of focal attention) but not others (e.g., detecting scene changes). What would be problematic for my proposal is a case in which there's visual experience of seen objects, but no visual information is available for voluntarily guiding attention.

5.4.2 Phenomenal Properties

The second objection specifically concerns the neurophysiological evidence presented in §5.3.3. The presented evidence, the objection goes, only supports that the higher regions of the ventral stream (IT cortex) play the attention-guiding role. It's a further step to say that visual experiences correlating with neural activity in these areas themselves play the attention-guiding role. And, even if the neurophysiological evidence did suggest that visual experiences play the attention-guiding role there would still be a problem. As Kriegel (2004) points out, showing that visual experiences play a given functional role is not enough to show that they do so in virtue of their experiential aspects. There is a difference between a visual experience and the *phenomenal properties* had by that experience. The phenomenal properties of a visual experience are those aspects or features of it in virtue of which it's phenomenally conscious at all (Chalmers 2004, 341, 2006, 50). Even if it's granted that visual experiences themselves fill the attention-guiding role it doesn't follow that their phenomenal properties have anything to do with it. It may be in virtue of other features of visual experience (e.g., simple neurophysiological features) that they manage to make visual information available for

use in the voluntary control of focal attention. The phenomenal properties of visual experiences may be idle with respect to this functional role.

The objection raises a serious issue, but is aimed at a strawman. Recall from §5.1 that the neurophysiological evidence (along with the introspective evidence) isn't being given as part of a deductive argument for the attention-guiding role of visual experience. The proposal that visual experience plays this role isn't supposed to follow logically from, or be entailed by, the evidence. Instead, the neurophysiological facts are presented here as interesting and compelling evidence that visual experience plays the attention-guiding role. Even acknowledging the gap between the operation of the neural correlates of visual experience and visual experience itself, and the further gap between visual experiences and their phenomenal properties, the neurophysiological facts do provide reason to think that visual experience plays the attention-guiding role.

A related point is that the objection amounts to an unreasonable demand. The objection uses the gap between neural correlates of visual experience and the experiential aspects (phenomenal properties) of visual experiences to undercut neurophysiological evidence for the functional role of visual experience. But the only way to respond to such an objection is to close the gap by giving some account of the metaphysics of visual experiences and their phenomenal properties. So, if the objection carried substantial weight, then that would mean that neurophysiological facts could be relevant to the role of visual experience only if this metaphysical account is already given. But surely that's too strong a demand—surely using neurophysiological results doesn't require first defending an account of the metaphysics of visual experiences.

5.5 Is Visual Experience Necessary?

Given its attention-guiding role, is visual experience necessary for vision-based thought? Or, is access consciousness to the information gathered through vision enough for vision-based thought?[19] More precisely, the question is:

> **Necessity Question:**
> Relative to some set of counterfactual circumstances, if you take away the experience in some instance of visually perceiving an object, but add back some other mode of access consciousness to the visual con-

---

[19]Smithies (2011b) also clearly separates the question of whether visual experience plays a role in vision-based thought from the question of whether visual experience is the only thing which could play that role.

> tent consistent with the counterfactual circumstances, would you re-
> store vision-based thought?

The necessity question has two parameters: (1) the set of counterfactual circumstances which provide the modal force of the necessity and (2) the mode of access consciousness being restored. The mode of access consciousness needs to be specified because there are a variety of mechanisms by which you could have access consciousness (e.g. Block 1995, 232). Some of these could allow for vision-based thought while others did not allow for it. This point already came out in §5.4.1 when it was noted that a mode of access consciousness might make information available for use in some cognitive functions but not others. The counterfactual circumstances need to be specified because they constrain the modes of access. For example, some modes of access might be metaphysically or conceptually possible, but not possible given human neurophysiological structure and cognitive functioning.[20]

Super blindsighters provide a standard way to frame the question and set the parameters (e.g., Campbell 2002). A super blindsighter, as Block imagines them (1995, 233), is a blindsighter who has access consciousness to visual content from her blindspot. While a normal blindsighter is only able to answer accurately questions about things in her blindspot when given a small set of choices from which to guess, a super blindsighter is able to prompt herself to make spontaneous guesses (i.e., without choices) about what's in her blindspot. The super blindsighter can do this in response to a variety of different needs, e.g. as a means to identify what's in the location of her blindspot or as a means to navigate around the environment. So the super blindsighter might simply wonder about what's in her blindspot and answer the question by making a spontaneous guess ("it's a red ball"). Or, she might spontaneously guess about whether and (if so) precisely how to turn to avoid an obstacle. We can ask then whether a super blindsighter, so conceived, could have vision-based thought about objects in her blindspot.

In this standard way of setting the parameters of the necessity question the mode of access consciousness is self-prompted spontaneous guessing (SSG) and

---

[20]In addition, some modes of access may not allow for vision-based thought given actual human neurophysiological structure and cognitive functioning, but may suffice in other circumstances. For example, on the proposal given here whether a given mode of access allows for vision-based thought depends on whether the way it affords access to visual information allows for its use in attentional control. But that, presumably, depends on details about how voluntary control of attention is facilitated, a factor up for specification within the counterfactual circumstances.

the counterfactual circumstances are (presumably) all metaphysically or conceptually possible circumstances. The usual intuition gotten by imagining what it would be like to be a super blindsighter is that, even allowing for all metaphysically or conceptually possible circumstances, having access to visual content from the blindspot through SSG would not allow for vision-based thought about objects in the blindspot (Campbell 2002; Smithies 2011b).

The proposal that visual experience plays the attention-guiding role provides two insights into the necessity question, both in the standard framing in terms of super blindsighters and more generally. The first insight is that since the functional role of visual experience is to make visual information available for use in the voluntary control of focal attention, restoring access consciousness will restore vision-based thought so long as the new mode of access consciousness still makes visual information available for use in the voluntary control of focal attention. So if the intuitions about super blindsighters noted above are correct, then SSG must not be a mode of access consciousness which makes visual information available for use in the voluntary control of attention. I think it's plausible that this correct: SSG is not a mode of access consciousness which makes visual information available in that way.

The second insight is that although adding SSG without visual experience might not enable vision-based thought, there might be other modes of access consciousness which when added do. This is because although SSG doesn't allow for the use of visual information in the voluntary control of attention, two empirical facts suggest other modes of access consciousness might allow for it. The first is that the FEF receive projections not just from the ventral stream, but also from the dorsal stream. If dorsal stream processing is left intact after a loss of visual experience, then there will still be a source of visual information for the FEF. The second fact is that the areas within the ventral stream itself are heavily interconnected; similarly, the ventral and dorsal streams are heavily interconnected (see Kravitz et al. 2013, 29–31 for review). These interconnections suggest that even when conscious seeing is not functioning normally information typically associated with it is still available to the FEF.

The upshot of the two facts is that the FEF still have access to a wealth of visual information even in the absence of visual experience. Of course, it's true that this potential availability doesn't allow for the voluntary control of attention in blindsight, the known case of nonconscious seeing. But that doesn't preclude the possibility of devising tricks or methods—alternatives to SSG—which enable

blindsighters to use effectively this visual information.[21] What's interesting or important is that there is the right sort of neural connectivity to the FEF to potentially allow for the voluntary control of attention.

Hence, although the standard framing of the necessity question in terms of super blindsighters and SSG might have a negative answer, if the proposal here is correct, then it might have a positive answer when alternative, empirically possible modes of access consciousness are considered. There might be empirically possible modes of access consciousness which, when added back, enable vision-based thought. Visual experience is not empirically necessary for vision-based thought.

---

[21]Before the forced-choice guessing procedure was developed blindsighters were unable to use their residual visual processing to identify features of objects within their blindfield. The suggestion here is that, similar to how forced-choice guessing enables blindsighters to use residual visual processing to identify features, other procedures might allow them to take advantage of the FEF connectivity to voluntarily control attention.

Chapter 6:
An Argument for Naïve Realism from Demonstrative Thought

6.1 The Problem of Hallucination

This concluding chapter has two aims. First, §6.2 extends the results from chapters 2 and 4 into an argument for naïve realism. Second, §6.3 collects together insights from the earlier chapters to present in outline a particular version of naïve realism—my view of experience. I suggest that this version of naïve realism is actually quite intuitive. This is important because, despite its name, naïve realism is often taken to be a radical, far-fetched view which goes against what both introspection and empirical psychology reveal about perceptual experience. As I discuss, this conception of naïve realism as radical and far-fetched is the result both of associating it with optional add-ons as well as an overly heavy reliance on intuitions about indiscriminability at the expense of other considerations.

   Naïve realism and representationalism are metaphysical views on the nature of perceptual experience *types*, the repeatable first-person conscious aspects instanced in particular cases of perceiving (see §2.4, page 37).[1] Naïve realism says perceptual experiences are relational states, while representationalism says that they are representational states (see §2.7).[2] More detailed characterizations will

---

[1]As the shape of the debate has changed over time, listing examples of each view is difficult. Broadly naïve realist views are held by Campbell (2002; 2010), Brewer (2006; 2007; 2011), Johnston (2006; 2011), Fish (2009), Leddington (2009), Kennedy (2011; 2013), Logue (2012a; 2012b), Hobson (2013), Knight (2013), and Genone (2014). Broadly representationalist views have been held by Burge (1991; 2005; 2009; 2010), Davies (1991; 1997), Tye (1995), Dretske (2003), Chalmers (2004; 2006), Byrne (2009), Pautz (2009; 2010), Speaks (2009), Schellenberg (2010; 2011; 2013), and Siegel (2010a; 2010b). Brewer (1999) has defended a representationalist view, but now rejects it. See also fn. 10 (page 14), which lists earlier examples of representationalism. Not all these examples developed their view in contrast to naïve realism, and some (e.g., Chalmers) wouldn't count as repesentationalists on the refined characterization of representationalism given in the next section (§6.2).

[2]Sometimes it's further specified by representationalists that experiences are propositional

114

be provided below in §6.2. but it will be helpful first to frame them as alternative ways to save direct realism from the problem of hallucination (see Crane 2006, 132–34; Genone forthcoming, who provide this framing; see also Dokic and Martin 2012 on this framing; see also Siegel 2006b, 356, 2010a, 176; Logue 2009, 18; Brewer 2011, 11 on direct realism).

The problem of hallucination takes the form of an argument which, using hallucinations, purports to show that you don't actually have perceptual experiences of the distal, mind-independent world. This is the denial of direct realism, called *indirect realism*. It has it that you perceive the world indirectly, through experiences which present you with mind-dependent objects. *Direct realism*, in contrast, is the view that in perceptual experiences you're presented with the distal, mind-independent world that you're perceiving (see §2.2, page 33; Genone forthcoming also provides a detailed discussion). There are many versions of this argument from hallucination (e.g., Ayer 1956, 90,95–113; Jackson 1977, 86,115; Valberg 1992a, ch 1; Robinson 1994; Smith 2002; Johnston 2004). The following version makes transparent how naïve realism and representationalism are competing ways to save direct realism (see Crane 2006, 133).

Argument from Hallucination

Premise 1: In an hallucinatory token experience there is no distal, mind-independent object that's perceived (definition of hallucination, see Chisholm 1957, 162; Fish 2010, 3).

Premise 2: Perceptual experience types are relations between the subjects who have the experiences and what's presented in them.

Premise 3: Any perceptual experience, i.e. an experience enjoyed while actually perceiving an object, could be instanced in a case of hallucinating.

Conclusion: What's presented in perceptual experiences are not the distal, mind-independent objects actually perceived.

To see that the conclusion follows from the premises, assume it's false. By premise

---

attitudes (e.g., Byrne 2009; Pautz 2010; see also French 2013). Chalmers (2004) provides a nuanced breakdown and discussion of representationalist views. Sometimes the distinction between naïve realism and repesentationalism is put in terms of typing token experiences by their phenomenal character (Logue 2009, 22-23; Siegel 2010a, 75–76; Nanay 2014b). In naïve realism the object perceived is relevant to this typing, while in repesentationalism its the content of the token experiences that matters. Although at first the two formulations seem different (Nanay 2014b), with the right distinctions they can be made equivalent.

2, it follows that perceptual experiences are relations to the distal, mind-independent objects perceived in them. Since relations are object dependent, instancing a perceptual experience requires that there actually is some distal, mind-independent object that's perceived. But clearly that contradicts premises 1 and 3: in hallucinations there's no perceived distal, mind-independent object (premise 1), but by premise 3 any perceptual experience can be instanced in a case of hallucination. Therefore, the conclusion cannot be false when all premises are true.

As Crane notes (2006), the above argument is a problem for direct realism because premise 1 is a definition and premises 2 and 3 are both compelling. In many versions of the argument from hallucination premise 3 is supported by the further premise that for any perceptual experience E there could be token hallucinatory experiences which instance an experience E′ that's indiscriminable from E (e.g., Valberg 1992a, 11–16). Premise 3 follows from this premise and the widely assumed principle that experiences indiscriminable to their subjects are identical.[3] Considerations about the transparency of experience are generally taken to make premise 2 an intuitive claim about experience. Experiences introspectively seem "transparent": when you introspect your experience all you find is what's presented. What's presented is immediately available to you. Premise 2 could be replaced by the weaker, but just as intuitive, premise that experiences are object dependent.[4]

Representationalism is a way to save direct realism by rejecting premise 2 and saving premise 3.[5] The representationalist denies that experience types are relations to what's presented in them, instead positing that they are representations. On one version, experiences are representations of what's presented in them.[6] Here the representationalist must lean on a preexisting characterization

---

[3]As noted in §1.4.1, Johnston (2011, 181) calls this principle the phenomenal bottleneck principle (see also Martin 2004, 40). The principle is widely believed, perhaps in part because it's often presented as definitional when the technical notion of experience (or phenomenal character) is being taught. It's explicitly endorsed by Schellenberg (2011, 738–40), who further derives it from assumptions about the introspectibility of changes in experience. Siegel (2010a, 169–70) and Chalmers (2006, 53–54) implicitly assume something like it; Chalmers (2006, 107–9) also explicitly endorses it.

[4]Fish (2009, 18–23) provides a discussion of transparency (see also Raleigh 2009; Millar 2014a,b). Valberg's assumption (1992a, 5) in his version of the argument that what's presented in experience is "temporally present" is an example. By this Valberg means that experiences are object dependent and he motivates the assumption through experience's transparency.

[5]Although not every representationalist view has this aim (e.g., Prosser 2011).

[6]Note all representationalists accept that experiences are representations of what's presented in them. For example, if you thought that experiences only had general content (Davies 1991; Tye 1995; Pautz 2009), i.e. denied that they represented particular objects, but still held that expe-

of representations, one which explains how an experience which presented a distal, mind-independent object could be instanced in an hallucination. The main option is to say that representational states (or representational properties, objects that are representations) are states that can be normatively evaluated with respect to some conditions in the world. They have accuracy or truth conditions (Crane 1992; Peacocke 1992a; Gunther 2003, 5; Burge 2009, 314, 2010, 292; Siegel 2010b).[7] Maps and assertions provide two examples (Dretske 1988, 59; see also Cummins and Poirier 2004, 22). A road map of Houston can be more or less accurate, e.g. you can ask whether the various streets it depicts are in the correct positions. Likewise, the assertion "In Houston, Westheimer runs perpendicular to Montrose Blvd." can be evaluated with respect to whether it gets the relation between the two named streets correct. What's important is that accuracy or truth-evaluable states are not necessarily object dependent. Maps can depict nonexistent cities and you can talk about what isn't there. So if experiences are representations of the distal, mind-independent objects presented in them, this potentially explains how they could be instanced in hallucinations (premise 3).

Naïve realism is a way to save direct realism by rejecting premise 3 and saving premise 2. Note that premise 2 isn't itself naïve realism, since it leaves open what's presented in an experience. Indirect realists accept premise 2, but unlike naïve realists reject that what's presented in experience are the distal, mind-independent objects perceived. Although naïve realists don't incur a burden to explain relations in the way representationalists must explain representations,[8] they do need to explain how hallucinations can instance experiences indiscriminable to perceptual experiences without simply identifying the two. Naïve real-

riences presented particular objects, then you'd need a more complicated connection between what's represented by and what's presented in experiences.

[7]A second option is to say that representational states (or properties, objects) are stand-ins for other states, properties, and objects (Dretske 1988, 52; Grush 1997, 5). To use an example from Dretske (1988, 52), I might explain a basketball play I saw to you by using miscellaneous nearby items (e.g., some coins in my pock and leftover popcorn). I move the items around as if they were the players during the play, using them to help explain what happened. The items would serve as stand-ins for, or representations of, the actual players.

[8]If told that being-to-the-left-of was a relation, it would be confused to ask, what relation *is* being-to-the-left-of? Similarily, it would be confused to ask a naïve realist, what relation *is* an experience? The claim is flatly that experiences are relations, not that there is some further, to-be-identified relation which constitutes them. Of course, it makes sense to ask whether the relation is sui generis (the standard naïve realist view, Fish 2009, 14; Logue 2009, 21) or reducible to other relations. It also makes sense to further ask about the higher-order properties of the relation (e.g., is it causal? spatial?). But none of these details need to be given before the claim that experiences are relations is a coherent claim which stands on its own.

ists cannot reject premise 3 out of hand and leave unexplained the possibility of indiscriminable hallucinations which motivated it.[9] Part of the work of §6.3 will be to give the best naïve realist explanation of indiscriminable hallucinations.

6.2   Extending the Argument to Naïve Realism

Recall that chapter 2 concluded that perceptual experiences are object dependent. Specifically, a perceptual experience of an object O is object dependent on O, in the sense that the experience can only be instanced in cases in which O is actually perceived (see §2.4, page 37). The chapter used two assumptions about perception-based thought (recall the main thesis, page 36), but here I'll take them for granted.[10] As pointed out in §2.7, the object dependence of experience is consistent with both naïve realism and representationalism. There are versions of representationalism on which perceptual experiences are object dependent (Genone forthcoming, §3). (These versions, of course, are inconsistent with adopting representationalism as a way to save direct realism from the problem of hallucination, since they make premise 3 false.) The strategy in this section will be to extend the result from chapter 4 to rule out representationalism, thereby leaving the argument from chapter 2 to support just naïve realism.[11] Chapter 4 did present its result as consistent with perceptual experiences being representations. The extension requires arguing that representationalism has commitments beyond the mere claim that experiences are representations. It is further committed to grounding the content of perceptual experiences in the content of representations constructed by the visual system.

---

[9]Similarly, representationalists cannot reject premise 2 out of hand. They need an explanation of the transparency of experience (see Fish 2009, 18–23; Millar 2014a,b for discussion).

[10]The second assumption that experience plays a role in perception-based thought is explicitly defended in chapter 5, at least for the case of vision. The first assumption, perception-based thought is demonstrative, I take to be widely assumed. The more important claim which followed from it, that perception-based thought is relational, is indirectly supported in chapter 3 by giving a plausible account of the exploited relation.

[11]Arguments for naïve realism tend to fall into a few kinds. Traditionally there are arguments from the transparency or particularity of experience (Martin 2006, 354–55; Fish 2009, 18–23; Logue 2012b; Hobson 2013, 555; Knight 2013, 3; although see Raleigh 2009; Sollberger 2012 for replies). Of course, arguments from experience's role in demonstrative thought are another variety (Campbell 2002, ch 6,7; Campbell 2004; Martin 2002a, 197–200). There are also Johnston's argument from experience's epistemic role, Martin's argument from sensory imagining, and Fish's argument from the hard problem of consciousness (Johnston 2006, 2011; Martin 2002b; Fish 2009, 75–79). Chapter 2 suggests that the arguments from demonstrative thought only show the object dependence of experience. The argument developed in this section is an original approach falling into none of kinds just listed.

Before giving the extension to rule out representationalism, it should be noted that ruling out representionalism isn't sufficient to establish naïve realism. For example, an indirect realist view such as sense-data theory might be correct (Russell 1997/1912; Price 1932; Ayer 1956; Jackson 1977; see also Locke 1997/1689; Hume 2000/1739–40). If so, then neither representationalism nor naïve realism are right. Like naïve realism, these indirect realist views take perceptual experiences to be relations (premise 2 of the argument from hallucination). Like representationalism they accept that hallucinations can instance perceptual experiences (premise 3). Unlike either naïve realism or representationalism, these indirect views accept the conclusion of the argument from hallucination. They hold that what's presented in experience is something besides the distal, mind-independent objects actually perceived—typically mind-dependent "sense data" or "ideas". The argument from chapter 2 rules out indirect realism, since experiences are not object dependent on indirect realist views. Arguing for naïve realism requires both the argument from chapter 2 and the extension of chapter 4: the argument from chapter 2 supports an object dependent view of experience, such as naïve realism or representationalism, while the extension of chapter 4 rules out the representationalist alternative.[12]

The extension of the argument from chapter 4 to rule out representationalism goes through three steps:

Step 1: Argue that if there's no representations of an object O, presented in a visual experience E, constructed in the visual system, then no visual representations of O are constructed in the brain.

Step 2: Next argue that if no visual representations of O are constructed in the brain, then no perceptual representations of O are constructed in the brain.

Step 3: Finally argue that if no perceptual representations of O are constructed in the brain, then representationalism is false.

Recall (§4.2.1, page 76) that the visual system is the functionally defined set of neurological processes which realize whatever (sub)tasks are involved in seeing. This makes step 1 straightforward. Even setting aside any stipulated definitions, a representation constructed in the brain will count as visual, presumably, only if it's involved in some subtask of seeing. Since the visual system is *defined* as

---

[12]There are yet other alternatives besides indirect realism, naïve realism, and representationalism. The main one is adverbialism (see Chisholm 1957), which denies that experiences are object directed or presentational at all. Presumably the argument from chapter 2 rules out this option as well, but the issue is less clear.

whatever processes realize seeing-related tasks (as opposed to, say, being defined biologically, e.g. through brain location or cytoarchitecture), then by definition any visual representations will be ones constructed in the visual system.

The argument needed for step 2 is less straightforward. A representation constructed in the brain will count as perceptual, presumably, only if it's involved in some perceptual processing task. With this in mind, step 2 turns on assuming that when it comes to describing the (functionally defined) subpersonal perceptual processing tasks which explain how you're able to perceive, these tasks will be exhaustively classifiable to one or more sensory modalities. Although a given perceptual processing task might belong to multiple sensory modalities, i.e. is involved in enabling perception in more than one sensory modality, no perceptual processing tasks will be amodal. That is, no perceptual processing tasks aren't involved in enabling perception in one sensory modality or another. So, perceptual representations are always constructed as part of some modality-associated perceptual processing task, i.e. are always constructed in some modality-specific perceptual system (e.g., the visual system or auditory system). Hence, if a perceptual representation for an object presented in a visual experience is constructed at all, it's constructed within some modality-specific perceptual system. Finally, it's implausible that no representation for an object presented in a visual experience would be constructed within the visual system, but one would be constructed within some nonvisual perceptual system.

Moving on to step 3, the first point to note is that experiences can be both relations and representations (Siegel 2010b, 363–365; Pautz 2011, 396; Logue 2014; see also French 2013). As noted in §6.1, whether or not something is a representation depends on whether it's the kind of thing which can be evaluated for truth or accuracy, i.e. the kind of thing which has content. The simple point is that even if experiences turn out to be relations as naïve realists propose, those relational states could still turn out to be truth or accuracy evaluable. For example, it might be that experiences are relations. But they're also relations with phenomenal, or first-person subjective aspects. Within visual experiences, for example, there's a way things "phenomenally look" to the subject, or appear to the subject based on the experiential aspects themselves (Chisholm 1957, ch 4; Jackson 1977, 30-41). These phenomenal aspects, or properties (see §5.4.2), of experiences might ground content (see Siewert 1998, ch 7; Horgan and Tienson 2002; see also Pautz 2013). Experiences might be representations in virtue of their phenomenology. It's worth noting that most representationalist arguments for the claim that experiences are representations take just this form: they argue

that phenomenology grounds content, and since (of course) naïve realists accept that experiences have phenomenal aspects, they should accept that experiences have content (e.g., Byrne 2009, 441–44; Pautz 2009, 48, 2011, 396; Siegel 2010a, 45, 2010b; Schellenberg 2011, 719–20; but not all do, see Nanay 2014a; see Nanay forthcoming for discussion). So, representationalism cannot be distinguished from naïve realism merely by the claim that experiences are representations.[13]

You might object in two ways. First, you might press that representations, by definition, are detachable from what's represented (Grush 1997). They're "intentional" states as Anscombe (1965) has described, in the sense that what's represented need not exist or be causally connected to them. Since this wouldn't be possible with relational states that are representations, relational states cannot be representations. But this objection ignores that there is, in fact, a well-defined notion of object dependent representation and its widely adopted by representationalists (see §2.7, page 46; for object dependent content, see Burge 1977, 346; Evans 1982; McDowell 1982, 204, 1984, 287; Martin 2002a, 178; Crane 2011, 23; see also Jeshion 2010b for recent work). It also seems to go conceptually beyond the basic idea of representations as states with truth or accuracy-evaluable content. Surely there's nothing about the idea of a state with content that entails that state cannot be dependent on what's represented. The second objection is that with representations there must be some possibility of inaccurate or false representation. But, the objection continues, experiences, being relations, could not misrepresent on the naïve realist view (and therefore can't be representations). In reply, the misrepresentation requirement itself could be rejected. There are many prima facie counterexamples, e.g. any representation the content of which is a necessary truth. But setting that point aside, there's no reason to think that experiences, if they were relations, would never misrepresent. What's true is that experiences could not *present* things not actually perceived. But whether this means they cannot represent inaccurately all depends on how their presentational or phenomenal aspects relate, or ground, their representational content.[14]

---

[13]It's true that many naïve realists reject this claim, mostly by trying to undercut the capacity of phenomenology to ground content (see Campbell 2002; Travis 2004; Fish 2009; Brewer 2011). For example, some press that phenomenology leaves the content too underdetermined (Travis 2004), while others argue that it would lead to inconsistent or impossible content (Brewer 2006, 170, 2011). But the basic point remains: there's nothing conceptual stopping relational states from also being representations.

[14]For example, perhaps different phenomenal features ground the same content. For the naïve realist, an experience of a blue ball under white light is different from the experience of a white ball under blue light. But if, e.g., discriminability plays a role in determining the content grounded

What is representationalism if not just the claim that experiences are representations? Clearly representationalists will want to reject the naïve realist picture on which experiences are sui generis relations to what's presented in them which happen, due to their phenomenal features, to have content. We can make progress on the commitments of representationalism by examining why representationalists reject this picture. There are two ways to put what goes wrong for the representationalist. First, on this picture experiences aren't *fundamentally* representations. Although it's common to use the term 'fundamental' to further refine both naïve realism and representationalism (e.g., Logue 2012a, 174; Schellenberg 2013, 49), precisely what's meant is often not clear. One thing that might be meant is that, in some recognizable sense, in the above naïve realist picture the experience's representational side plays no role in explaining what it is for a state to be an experience. All the interesting work was in articulating how experiences are relations—that these relational states also have content is a coincidental fact noticed later. Second, the content of experience isn't playing the right role. For example, at least for some representationalists having a certain content should explain the experience's phenomenology (Tye 1995; Dretske 2003; Pautz 2009; Speaks 2009). If the phenomenology is a matter of being in a relation to what's perceived, and the phenomenology grounds content, then this reverses the explanatory roles. On the naïve realist picture it's the phenomenology which explains the content.[15]

These considerations suggest that representionalists are committed to a certain account or range of accounts of what grounds the content of experiences.[16] The representationalist account should ground content in features of experiences which still allow that content to play a substantial role in explaining the nature

---

by given phenomenal features, then these experiences might have the same content.

[15]It's not only naïve realists who suggest that perceptual content is grounded in perceptual phenomenology (e.g., see Horgan and Tienson 2002). Others have also pointed out that the explanatory role of content is an important point in the dispute between naïve realists and representationalists (see Chalmers 2004; Genone forthcoming).

[16]Note that this suggestion holds even if the representationalist holds that nothing grounds the content of experiences. Experiences, so this view goes, are sui generis represenional states. This representationalist will still have to explain what makes a given *token* experience an instance of one experience type over another. Even if experience types are sui generis representational states, some story needs to be given of what makes one token experience an instance of one type over another. Such an explanation will amount to what I'm calling an account of how content is grounded, and certain versions will be ruled out for the considerations given above. For example, this representationalist wouldn't want to say that what makes a token experience an instance of one type over another is its phenomenology.

of experiences. One possible account would ground the content of experiences in the content of the perceptual representations in the brain which bring about the experiences. For example, an experience of O might constitutively involve the construction of a perceptual representation of O in the brain; having an experience of O would then be a representational state the content of which was grounded in the underlying constitutive perceptual representations. This kind of account would leave room for explaining phenomenology in terms of representation. Unfortunately, the results of chapter 4 along with steps 1 and 2 just completed above rule it out.

To complete step 3, the suggestion now is that accounts which ground the content of experience in perceptual representations are the only ones which give content the right role or make experiences fundamentally representations. Accounts which ground content in phenomenology clearly won't work. What are the other options? The only other option that's apparent to me is to appeal to perception-related relations that hold between what's represented and the experience itself. A natural candidate, for example, is causal covariation: perhaps a given perceptual experience causally covaries with a given perceived object or property, and it's that relation of covariation which grounds the content. Any such account will likely be tied to an already existing account of representation, e.g. Fodor's asymmetric dependency view or Millikan's teleological biosemantics (Fodor 1987; Millikan 2004).

There is a major obstacle to these views. As Siegel points out (2010a, 43), for an experience to have content requires that the content is in some way conveyed to, or accessible by, the subject who has the experience. One thing that makes the accounts of grounding in terms of phenomenology and perceptual representations attractive is that they face no prima facie challenges on this front. Although details are still needed, it should seem plausible that content grounded by phenomenology is accessible to the subject. Likewise, perceptual representations are the kinds of things which are integrated into your overall cognitive architecture—and thereby are well poised to ground content that's accessible to you. But pending the details it's unclear why content grounded in (say) causal relations to the world would be accessible to you. Thinking about the issue in informational terms, most of the information available in perceptual experiences is not accessible. Information that is available is information that's tracked or used by perceptual systems as they construct representations.

The above considerations are far from definitive, but nonetheless push the burden onto the representationalist. Representationalism requires a certain type

of account of how experiential content is grounded, but with grounding in terms of perceptual representations in the brain off the table there are no good candidates around.

6.3   A Sensible View of Experience

What are perceptual experiences? A particular picture of them starts to emerge from the proceeding work (chapters 2–6.2).

The View of Experiences:

Object-Dependency: What's perceived in an experience is constitutive of it.

Processing-Dependency: The construction of perceptual representations in the brain is constitutive of experiences, but not necessarily representations of what's experienced.

Phenomenology-Grounding: The phenomenal character of a token experience (the experience it instances) is determined by, or grounded in, these two constitutive components.

Content-Grounding: The representational content of an experience is determined by, or grounded in, its phenomenal character.

Anti-Disjunctivism: Exactly this account applies to both perceptual and hallucinatory experiences.

In the remainder of this section I will fill out the rudiments of this outline.

First, note that the view is thoroughly naïve realist without any compromise towards representationalism. Experiences, on this view, are relational states that involve (object-dependency) the distal, mind independent things perceived as well as (processing-dependency) neural processing of causally-derived information from those things. It's these two factors which determine the identity of the experience (phenomenology-grounding). The view allows that these states are representations, but effectively deprives the representational content of any substantial explanatory role (content-grounding). Experiences are not fundamentally representations, even if they have representational content.

Next consider the motivations behind the view. The object-dependency thesis is, of course, motivated by the considerations from chapter 2. The processing-dependency thesis is motivated first by run-of-the-mill facts about perceiving. Even rudimentary observation of the operation of the brain during perceiving

and affects of neurobiolological changes on experiences suggests that this processing is constitutive of experience, while the argument from chapter 4 suggests against a one-to-one match between representations constructed in that processing and what's presented in experience. The content-grounding thesis falls out of the considerations of the previous section, §6.2.

The above overview of motivations shows that the view is a sensible view of experience. It's what falls out when you take seriously (1) the functional role of perceptual experience (e.g., making objects available for thought), along with (2) the best contemporary perceptual psychology, and (3) arguments that phenomenology grounds representational content (e.g., Siewert 1998, ch 7; Horgan and Tienson 2002; Byrne 2009, 441–44; Pautz 2009, 48, 2011, 396, 2013; Siegel 2010a, 45, 2010b; Schellenberg 2011, 719–20). In other words, the naïve realist view just outlined is what naturally falls out of a broad approach to studying the nature of experience.

As noted in §6.1, this is in contrast to the usual picture of naïve realism as a radical, far-fetched view. The blame for this bad reputation rests both with naïve realists and their opponents, representationalists. On the naïve realist side, many naïve realists replace the content-grounding thesis with a more radical anti-representation thesis which denies that experiences have content (e.g., Campbell 2002; Travis 2004; Fish 2009; Brewer 2011). While there might be something to say for some of the arguments against experiences having representational content, a flat denial of experiences being representational is, on the whole, implausible. Also on the naïve realist side is a tendency towards disjunctivism.[17] For a long time the thought was that if experiences enjoyed while perceiving are relations to the distal, mind-independent things perceived, since then in hallucinations there is no distal, mind-independent object perceived, no experience of a kind enjoyed while perceiving could be instanced. In other words, hallucinations and perceptual experiences cannot be the same type of psychological state. For various reasons, disjunctivism is a deeply problematic view. Accepting it only further radicalizes naïve realism.

On the anti-naïve realist, representationalist side, naïve realism looks radical because of an undue focus on a single aspect of experience (as opposed to the broad approach canvassed here). Representationalists have tended to be impressed by the fact that hallucinations can instance experiences which are indis-

---

[17]Disjunctivism itself is out of the bounds of this discussion. Byrne and Logue (2008) provide a useful introduction. See also the collections from Haddock and Macpherson (2008) and Byrne and Logue (2009).

criminable to perceptual experiences. For example, right now I'm looking at a #2 pencil, having a veridical perceptual experience. But, for all I can tell through introspection, I might be hallucinating. In that case, I'd be instancing an experience which was at minimum indiscriminable from the actual perceptual experience I'm now enjoying. It's tempting to take the further step and say that this indiscriminable (hallucinatory) experience would be identical to the actual enjoyed perceptual experience. Part of this, perhaps, is a further deep commitment to the idea that experiential, i.e. phenomenal, aspects of perceiving must be fully accessible to the subject through introspection. This is Johnston's bottleneck principle (see fn. 3, page 116).

Although the bottleneck principle and its accompanying explanation of indiscriminable hallucinations is tempting, what I'm suggesting is that once the big picture is in view it loses much of its appeal. Once you consider things like the functional role of perceptual experience and see alternative ways to explain indscriminability, it becomes reasonable to reject the bottleneck principle.[18] This leads to the final issue: just how does the naïve realist view outlined here explain indiscriminability and, relatedly, why isn't disjunctivism obligatory? The central insight here comes from Heather Logue (2012b) and turns on the processing-dependency thesis. Naïve realists, as Logue points out, aren't limited to appealing to the perceived distal, mind-independent objects and their perceived properties to explain experiences (to explain the experiential aspects to a given instance of perceiving). They can also also to the subject end of the relation—to the neural processing within the head that's also constitutive of experiences. Take, for example, the experience I enjoy when actually a moment ago when looking at a #2 pencil and the potential indiscriminable experience I could enjoy were I only hallucinating. These two experiences are different because certain constitutive com-

---

[18]In addition, the bottleneck principle faces a serious challenge (see Johnston 2011). Assume the principle is true. Then any two introspectively indiscriminable experiences are the same. Now, imagine looking at a series of color patches each of which is slightly different from the next, so that the difference is color is so small that they look the same to you, but overall the first and last patches are obviously different. So, for any two pairs of color patches in this series, when you view them you have introspectively indiscriminable experiences. By the bottleneck principle, for any pair of adjacent patches in the series you have the same experience when viewing them. Since identity is transitive, you must be having the same experience no matter which color patches in the series you're viewing. But you're obviously not: since the first and last patches look visibly different to you, the experiences you enjoy while looking at them are different. The conclusion is that the bottleneck principle is wrong. Of course, there are ways to respond to this argument. The point is that the bottleneck principle, however initially plausible, is far from being free of difficulties on its own.

ponents in the perceptual case, e.g. the pencil, are missing from the hallucination. But other constitutive components are presumably the same: presumably the underlying neural processing in both cases would be the same. This sameness of neural processing between the perceptual and hallucinatory cases also explains their indiscriminability. It's plausible that only differences in the experiential aspects of instances of perceiving which are due to differences in underlying neural processing are differences which are discernible through introspection.

Bibliography

Akins, Kathleen. 1996. "Of Sensory Systems and the "Aboutness" of Mental States". *The Journal of Philosophy* 93:337–72. 72

Alais, David, Blake, Randolph, and Lee, Sang-Hun. 1998. "Visual features that vary together over time group together over space". *Nature Neuroscience* 1:160–64. 87

Allport, Alan. 2011. "Attention and Integration". In Christopher Mole, Declan Smithies, and Wayne Wu (eds.), *Attention: Philosophical & Psychological Essays*. Oxford University Press, 24–59. 40, 57, 58, 87

Anscombe, G. E. M. 1965. "The Intentionality of Sensation: A Grammatical Feature". In Ronald J. Butler (ed.), *Analytic Philosophy*. Blackwell, 158–80. 121

Armstrong, Katherine M. 2011. "Covert Spatial Attention and Saccade Planning". In Christopher Mole, Declan Smithies, and Wayne Wu (eds.), *Attention: Philosophical & Psychological Essays*. Oxford University Press, 78–96. 58, 104

Avrahami, Judith. 1999. "Objects of attention, objects of perception". *Perception & Psychophysics* 61:1604–1612. 87

Ayer, A. J. 1956. *The Problem of Knowledge: An Enquiry into the Main Philosophical Problems that enter into the Theory of Knowledge*, chapter 3. Pelican Books. 26, 115, 119

Baars, Bernard J. 1988. *A Cognitive Theory of Consciousness*. Cambridge University Press. 94

Bach, Kent. 1987. *Thought and Reference*. Oxford University Press. 11, 44, 56, 68

———. 2010. "Getting a Thing into a Thought". In Robin Jeshion (ed.), *New Essays on Singular Thought*. Oxford University Press, 39–63. 2, 6, 49, 53, 55, 93

Barcan Marcus, Ruth. 1961. "Modalities and Intensional Languages". *Synthese*

13:303–322. 53

Beck, Jacob. 2012. "The Generality Constraint and the Structure of Thought". *Mind* 121:563–600. 90

Behrmann, Marlene, Zemel, Richard S., and Mozer, Michael C. 1998. "Object-based attention and occlusion: Evidence from normal participants and a computational model". *Journal of Experimental Psychology* 24:1011–1036. 87

Block, Ned. 1995. "On a confusion about a function of consciousness". *Behavioral and Brain Sciences* 18:227–287. 16, 17, 18, 34, 61, 73, 97, 108, 109, 111

——. 2010. "Attention and Mental Paint". *Philosophical Issues* 20:23–63. 26

Bourget, David and Mendelovici, Angela. 2014. "Tracking Representationalism: Lycan, Dretske, and Tye". In Andrew Bailey (ed.), *Philosophy of Mind: The Key Thinkers*. Continuum, 209–35. 3

Brewer, Bill. 1999. *Perception and Reason.* Oxford University Press. 114

——. 2006. "Perception and Content". *European Journal of Philosophy* 14:165–181. 3, 22, 25, 26, 30, 34, 114, 121

——. 2007. "Perception and its Objects". *Philosophical Studies* 132:87–97. 25, 114

——. 2011. *Perception and its Objects.* Oxford University Press. 3, 14, 25, 26, 31, 34, 46, 49, 74, 114, 115, 121, 125

Brogaard, Berit. 2011a. "Are There Unconscious Perceptual Processes?" *Consciousness and Cognition* 20:449–63. 16, 96

——. 2011b. "Conscious Vision for Action Versus Unconscious Vision for Action". *Cognitive Science* 35:1076–1104. 94, 104

——. 2012. "Non-Visual Consciousness and Visual Images in Blindsight". *Consciousness and Cognition* 21:595–596. 16, 96

Brooks, Rodney. 1991. "Intelligence without representation". *Artificial Intelligence* 47:139–59. 12, 72

Bruce, Charles J. and Goldberg, Michael E. 1985. "Primate Frontal Eye Fields. I. Single Neurons Discharging Before Saccades". *Journal of Neurophysiology* 53:603–635. 103, 104

Bruce, Charles J., Goldberg, Michael E., Bushnell, M. Catherine, and Stanton, Gregory B. 1985. "Primate Frontal Eye Fields. II. Physiological and Anatomical Correlates of Electrically Evoked Eye Movements". *Journal of Neurophysiology* 54:714–734. 103

Buffalo, Elizabeth A., Bertini, Giuseppe, Ungerleider, Leslie G., and Desimone, Robert. 2005. "Impaired Filtering of Distracter Stimuli by TE Neurons following V4 and TEO Lesion in Macaques". *Cerebral Cortex* 15:141–151. 107

Burge, Tyler. 1977. "Belief De Re". *The Journal of Philosophy* 74:338–62. 27, 121

———. 1991. "Vision and Intentional Content". In Ernest Lepore and Robert van Gulick (eds.), *John Searle and His Critics*. Blackwell, 195–214. 25, 114

———. 2005. "Disjunctivism and Perceptual Psychology". *Philosophical Topics* 33:1–78. 3, 14, 22, 23, 25, 27, 35, 37, 74, 114

———. 2009. "Perceptual Objectivity". *Philosophical Review* 118:285–324. 25, 114, 117

———. 2010. *Origins of Objectivity*. Oxford University Press. 25, 27, 47, 74, 76, 114, 117

Byrne, Alex. 2009. "Experience and Content". *Philosophical Quarterly* 59:429–451. 14, 25, 37, 114, 115, 121, 125

Byrne, Alex and Logue, Heather. 2008. "Either/Or". In Adrian Haddock and Fiona Macpherson (eds.), *Disjunctivism: Perception, Action, Knowledge*. Oxford University Press, 57–94. 125

Byrne, Alex and Logue, Heather (eds.). 2009. *Disjunctivism: Contemporary Readings*. The MIT Press. 125

Campbell, John. 1997. "Sense, Reference, and Selective Attention". *Proceedings of the Aristotelian Society, Supplementary Volume* 71:55–74. 2, 18, 51, 56, 57, 61, 94, 98

———. 2002. *Reference and Consciousness*. Oxford University Press. 2, 3, 14, 17, 18, 19, 20, 22, 25, 26, 30, 31, 32, 34, 35, 37, 40, 46, 49, 50, 53, 54, 56, 57, 59, 61, 74, 93, 94, 95, 98, 111, 112, 114, 118, 121, 125

———. 2004. "Reference as Attention". *Philosophical Studies* 120, Proceedings of the Thirty-Fifth Oberlin Colloquium in Philosophy: Philosophy of Percep-

tion:265–276. 18, 19, 20, 30, 56, 94, 98, 118

———. 2010. "Demonstrative Reference, the Relational View of Experience, and the Proximality Principle". In Robin Jeshion (ed.), *New Essays on Singular Thought*. Oxford University Press, 193–212. 3, 22, 25, 34, 114

———. 2011a. "Relational vs Kantian Responses to Berkeley's Puzzle". In Johannes Roessler, Hemdat Lerman, and Naomi Eilan (eds.), *Perception, Causation, and Objectivity*. Oxford University Press, 35–50. 19

———. 2011b. "Visual Attention and the Epistemic Role of Consciousness". In Christopher Mole, Declan Smithies, and Wayne Wu (eds.), *Attention: Philosophical & Psychological Essays*. Oxford University Press, 323–41. 18, 20

Carrasco, Marisa. 2011. "Visual attention: The past 25 years". *Vision Research* 51:1484–1525. 57, 58, 59, 88

Carrasco, Marisa, Ling, Sam, and Read, Sarah. 2004. "Attention alters appearance". *Nature Neuroscience* 7:308–13. 57

Cassam, Quassim. 2011. "Tackling Berkeley's Puzzle". In Johannes Roessler, Hemdat Lerman, , and Naomi Eilan (eds.), *Perception, Causation, and Objectivity*. Oxford University Press, 18–34. 19

Chalmers, David J. 2000. "What is a Neural Correlate of Consciousness?" In Thomas Metzinger (ed.), *Neural Correlates of Consciousness*. MIT Press. Page citations to the reprint in David Chalmers. 2010. *The Character of Consciousness*. Oxford University Press. 74, 103

———. 2004. "The Representational Character of Experience". In Brian Leiter (ed.), *The Future for Philosophy*. Oxford University Press, 153–181. Page citations to the reprint in David Chalmers. 2010. *The Character of Consciousness*. Oxford University Press. 25, 37, 47, 109, 114, 115, 122

———. 2006. "Perception and the Fall From Eden". In Tamar Szabo Gendler and John Hawthorne (eds.), *Perceptual Experience*. Oxford University Press, 49–125. 21, 25, 29, 37, 47, 109, 114, 116

Chirimuuta, Mazviita and Gold, Ian. 2009. "The Embedded Neuron, the Enactive Field?" In John Bickle (ed.), *The Oxford Handbook of Philosophy and Neuroscience*. Oxford University Press, 200–225. 77

Chisholm, Roderick. 1957. *Perceiving: a Philosophical Study*. Cornell University

Press. 115, 119, 120

Clark, Austen. 2004. "Feature-Placing and Proto-Objects". *Philosophical Psychology* 17:443–469. 77

——. 2006. "Attention & Inscrutability: A Commentary on John Campbell, "Reference and Consciousness" for the Pacific APA Meeting, Pasadena, California, 2004". *Philosophical Studies* 127:167–193. 19, 40, 59

Cohen, Jeremiah Y., Heitz, Richard P., Schall, Jeffrey D., and Woodman, Geoffrey F. 2009a. "On the Origin of Event-Related Potentials Indexing Covert Attentional Selection During Visual Search". *Journal of Neurophysiology* 102:2375–2386. 104

Cohen, Jeremiah Y., Heitz, Richard P., Woodman, Geoffrey F., and Schall, Jeffrey D. 2009b. "Neural Basis of the Set-Size Effect in Frontal Eye Field: Timing of Attention During Visual Search". *Journal of Neurophysiology* 101:1699–1704. 104

Crane, Tim. 1992. "The Nonconceptual Content of Experience". In Tim Crane (ed.), *The Contents of Experience: Essays on Perception*. Cambridge University Press, 136–157. 117

——. 2006. "Is There a Perceptual Relation?" In Tamar Szabo Gendler and John Hawthorne (eds.), *Perceptual Experience*. Oxford University Press, 126–146. 3, 46, 115, 116

——. 2011. "The Singularity of Singular Thought". *Aristotelian Society Supplementary Volume* 85:21–43. 27, 121

Crick, Francis and Koch, Christof. 1995. "Are we aware of neural activity in primary visual cortex?" *Nature* 375:121–23. 105

Cummins, Robert C. and Poirier, Pierre. 2004. "Representation and Indication". In Hugh Clapin, Phillip Staines, and Peter Slezak (eds.), *Representation in Mind: New Approaches to Mental Representation*. Elsevier, 21–40. 6, 117

Datta, Ritobrato and DeYoe, Edgar A. 2009. "I know where you are secretly attending! The topography of human visual attention revealed with fMRI". *Vision Research* 49:1037–1044. 63

Davies, Martin. 1991. "Individualism and Perceptual Content". *Mind* 100:461–84. 15, 25, 27, 47, 114, 116

——. 1992. "Perceptual Content and Local Supervenience". *Proceedings of the Aristotelian Society* 66:21–45. 14, 47

——. 1997. "Externalism and Experience". In Ned Block, Owen Flanagan, and Güven Guzeldere (eds.), *The Nature of Consciousness*. The MIT Press, 309–327. 25, 114

Dennett, Daniel. 1978. "Toward a cognitive theory of consciousness". *Minnesota Studies in the Philosophy of Science* 9:201–28. 74

Desimone, Robert and Duncan, John. 1995. "Neural Mechanisms of Selective Visual Attention". *Annual Review of Neuroscience* 18:193–222. 58

Dickie, Imogen. 2010. "We are Acquainted with Ordinary Things". In Robin Jeshion (ed.), *New Essays on Singular Thought*. Oxford University Press, 213–245. 2, 11, 31, 41, 44, 49, 51, 53, 55, 61, 63, 67, 68, 80, 91, 93, 94, 98

——. 2011. "Visual Attention Fixes Demonstrative Reference by Eliminating Referential Luck". In Christopher Mole, Declan Smithies, and Wayne Wu (eds.), *Attention: Philosophical & Psychological Essays*. Oxford University Press, 292–322. 11, 32, 44, 49, 51, 55, 56, 61, 66, 68, 91, 93, 95, 98

——. forthcoming. "Perception and Demonstratives". In Mohan Matthen (ed.), *Oxford Handbook of Philosophy of Perception*. Oxford University Press, –. 2, 49, 51, 52, 53, 55

Dokic, Jérôme and Martin, Jean-Rémy. 2012. "Disjunctivism, Hallucination and Metacognition". *WIREs Cognitive Science* 3:533–543. 115

Drayson, Zoe. 2012. "The Uses and Abuses of the Personal/Subpersonal Distinction". *Philosophical Perspectives* 26:1–18. 76

Dretske, Fred. 1981. *Knowledge and the Flow of Information*. The MIT Press. 3, 50

——. 1988. *Explaining Behavior: Reasons in a World of Causes*. The MIT Press. 3, 6, 50, 117

——. 1995. *Naturalizing the Mind*. The MIT Press. 3, 50

——. 2003. "Experience as Representation". *Philosophical Issues* 13:67–82. 3, 25, 47, 74, 114, 122

Eckardt, Barbara Von. 2012. "The representational theory of mind". In Keith

Frankish and William M. Ramsey (eds.), *The Cambridge Handbook of Cognitive Science*. Cambridge University Press, 29–49. 12

Evans, Gareth. 1973. "The Causal Theory of Names". *Aristotelian Society: Supplementary Volume* 47:187–208. 11, 67

——. 1982. *Varieties of Reference*. Oxford University Press. 5, 7, 11, 14, 27, 31, 53, 55, 67, 121

Feldman, Jacob. 1999. "The role of objects in perceptual grouping". *Acta Psychologica* 102:137–163. 87

Fish, William. 2009. *Perception, Hallucination, and Illusion*. Oxford University Press. 14, 25, 26, 33, 37, 48, 114, 116, 117, 118, 121, 125

——. 2010. *Philosophy of Perception: A Contemporary Introduction*. Routledge Contemporary Introductions to Philosophy. Routledge. 115

Fisher, Justin C. 2007. "Why Nothing Mental is Just in the Head". *Noûs* 41:318–334. 21

Fodor, Jerry A. 1975. *The Language of Thought*. Harvard University Press. 5, 55

——. 1987. *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. The MIT Press. 3, 50, 123

——. 1990. "A Theory of Content II: The Theory". In *A Theory of Content*. The MIT Press, 89–136. 3

——. 1998. *Concepts: Where Cognitive Science Went Wrong*. Oxford University Press. 7, 33, 55

——. 2008. *LOT 2: The Language of Thought Revisited*. Oxford University Press. 5, 55

Fodor, Jerry A. and Pylyshyn, Zenon W. 1981. "How Direct Is Visual Perception?: Some Reflections on Gibson's 'Ecological Approach'". *Cognition* 9:139–196. 76

French, Craig. 2013. "Perceptual Experience and Seeing That P". *Synthese* 190:1735–1751. 115, 120

Genone, James. 2014. "Evidential Constraints on Singular Thought". *Mind and Language* 29:1–25. 25, 56, 114

———. forthcoming. "Recent Work on Naive Realism". *American Philosophical Quarterly* –. 25, 46, 115, 118, 122

Gertler, Brie. 2012. "Understanding the Internalism-Externalism Debate: what is the Boundary of the Thinker?" *Philosophical Perspectives* 26:51–75. 25

Gibson, James J. 1966. *The Senses Considered as a Perceptual System*. Houghton Mifflin Company. 12, 72

———. 1986. *The Ecological Approach to Visual Perception*. Lawrence Erlbaum. 12, 72

Goodale, M.A. and Milner, A.D. 1992. "Separate Visual Pathways for Perception and Action". *Trends in Neurosciences* 15:20–25. 104

Gregory, Richard. 1970. *The Intelligent Eye*. Weidenfeld & Nicolson. 79

Grush, Rick. 1997. "The Architecture of Representation". *Philosophical Psychology* 10:5–23. 117, 121

Gunther, York (ed.). 2003. *Essays on Nonconceptual Content*. The MIT Press. 117

Haddock, Adrian and Macpherson, Fiona (eds.). 2008. *Disjunctivism: Perception, Action, Knowledge*. Oxford University Press. 125

Harman, Gilbert. 1990. "The Intrinsic Quality of Experience". In James Tomberlin (ed.), *Philosophical Perspectives*, volume 4. Ridgeview, 31–52. 14

Hawthorne, John and Scala, Mark. 2000. "Seeing and Demonstration". *Philosophy and Phenomenological Research* 61:199–206. 31, 49

Heitz, Richard P., Cohen, Jeremiah Y., Woodman, Geoffrey F., and Schall, Jeffrey D. 2010. "Neural Correlates of Correct and Errant Attentional Selection Revealed Through N2pc and Frontal Eye Field Activity". *Journal of Neurophysiology* 104:2433–2441. 104

Hobson, Kenneth. 2013. "In Defense of Relational Direct Realism". *European Journal of Philosophy* 21:550–574. 25, 46, 74, 114, 118

Hoffman, Donald and Richards, Whitman. 1984. "Parts of Recognition". *Cognition* 18:65–96. 83, 84

Horgan, Terence E. and Tienson, John L. 2002. "The Intentionality of Phenomenology and the Phenomenology of Intentionality". In David Chalmers

(ed.), *Philosophy of Mind: Classical and Contemporary Readings*. Oxford University Press, 520–533. 120, 122, 125

Hubel, David H. and Wiesel, Torsten N. 1959. "Receptive Fields of Single Neurones in the Cat's Striate Cortex". *The Journal of physiology* 148:574–591. 77

——. 1962. "Receptive Fields, Binocular Interaction and Functional Architecture in the Cat's Visual Cortex". *The Journal of physiology* 160:106–54. 77

——. 1968. "Receptive fields and functional architecture of monkey striate cortex". *Journal of Physiology* 195:215–243. 77

Hume, David. 1993/1777. *An Enquiry Concerning Human Understanding*. Hackett Publishing Company, 2 edition. With A Letter from a Gentleman to His Friend in Edinburgh and Hume's Abstract of a Treatise of Human Nature. 3

——. 2000/1739–40. *A Treatise of Human Nature*. Oxford Philosophical Texts. Oxford University Press. Edited by David F. Norton and Mary J. Norton. 119

Hutchinson, J. Benjamin and Turk-Browne, Nicholas B. 2012. "Memory-Guided Attention: Control From Multiple Memory Systems". *Trends in Cognitive Sciences* 16:576–579. 90

Hutto, Daniel D. and Myin, Erik. 2013. *Radicalizing Enactivism: Basic Minds without Content*. The MIT Press. 12, 72

Jackson, Frank. 1977. *Perception: A Representative Theory*. Cambridge University Press. 26, 115, 119, 120

Jensen, Melinda S., Yao, Richard, Street, Whitney N., and Simons, Daniel J. 2011. "Change Blindness and InInattention Blindness". *WIREs Cognitive Science* 2:529–546. 16, 109

Jeshion, Robin. 2010a. "Introduction". In Robin Jeshion (ed.), *New Essays on Singular Thought*. Oxford University Press, 1–36. 52, 54, 56, 61, 94, 98

Jeshion, Robin (ed.). 2010b. *New Essays on Singular Thought*. Oxford University Press. 53, 121

Jeshion, Robin. 2010c. "Singular Thought: Acquaintance, Semantic Instrumentalism, and Cognitivism". In Robin Jeshion (ed.), *New Essays on Singular Thought*. Oxford University Press, 105–140. 2, 11, 31, 44, 49, 53, 54, 55, 56, 68, 93

Johnston, Mark. 2004. "The Obscure Object of Hallucination". *Philosophical Studies* 120:113–83. 115

———. 2006. "Better Than Mere Knowledge? The Function of Sensory Awareness". In Tamar Szabo Gendler and John Hawthorne (eds.), *Perceptual Experience*. Oxford University Press, 260–290. 3, 17, 22, 25, 30, 31, 32, 34, 49, 53, 95, 98, 114, 118

———. 2011. "On a Neglected Epistemic Virtue". *Philosophical Issues* 21:165–218. 6, 12, 14, 21, 25, 31, 37, 46, 49, 74, 114, 116, 118, 126

Kahneman, Daniel, Treisman, Anne, and Gibbs, Brian. 1992. "The Reviewing of Object Files: Object-Specific Integration of Information". *Cognitive Psychology* 24:175–219. 20, 77, 78, 79, 83

Kelly, Sean D. 2004. "Reference and Attention: A Difficult Connection". *Philosophical Studies* 120:277–86. 18, 32, 34, 49, 95, 98

Kennedy, Matthew. 2011. "Naïve Realism, Privileged Access, and Epistemic Safety". *Noûs* 45:77–102. 25, 114

———. 2013. "Explanation in Good and Bad Experiential Cases". In Fiona Macpherson and Dimitris Platchias (eds.), *Hallucination: Philosophy and Psychology*. The MIT Press, 221–54. 25, 114

Kentridge, Robert W. 2011. "Attention Without Awareness: A Brief Review". In Christopher Mole, Declan Smithies, and Wayne Wu (eds.), *Attention: Philosophical & Psychological Essays*. Oxford University Press, 228–246. 91, 92

Kentridge, Robert W., Nijboer, Tanja C. W., and Heywood, Charles A. 2008. "Attended but unseen: Visual attention is not sufficient for visual awareness". *Neuropsychologia* 46:864–69. 91

Knight, Gordon. 2013. "Disjunctivism Unmotivated". *Phenomenology and the Cognitive Sciences* 1–18. 25, 46, 74, 114, 118

Kravitz, Dwight J. and Behrmann, Marlene. 2011. "Space-, object-, and feature-based attention interact to organize visual Scenes". *Attention, Perception, & Psychophysics* 73:2434–2447. 64

Kravitz, Dwight J., Saleem, Kadharbatcha S., Baker, Chris I., and Mishkin, Mortimer. 2011. "A new neural framework for visuospatial processing". *Nature Reviews Neuroscience* 12:217–230. 104

Kravitz, Dwight J., Saleem, Kadharbatcha S., Baker, Chris I., Ungerleider, Leslie G., and Mishkin, Mortimer. 2013. "The ventral visual pathway: an expanded neural framework for the processing of object quality". *Trends in Cognitive Sciences* 17:26–49. 104, 105, 106, 112

Kriegel, Uriah. 2004. "The functional role of consciousness: A phenomenological approach". *Phenomenology and the Cognitive Sciences* 3:171–193. 94, 95, 109

Kripke, Saul. 1981. *Naming and Necessity*. Oxford: Blackwell. 53

Leddington, Jason. 2009. "Perceptual Presence". *Pacific Philosophical Quarterly* 90:482–502. 25, 114

Lee, Kyoung-Min, Ahn, Kyung-Ha, and Keller, Edward L. 2012. "Saccade Generation by the Frontal Eye Fields in Rhesus Monkeys Is Separable from Visual Detection and Bottom-Up Attention Shift". *PLoS ONE* 7:e39886. 104

Levine, J. 2008. "Demonstrative Concepts". *Croation Journal of Philosophy* 8:328–336. 33

Levine, Joseph. 2010. "Demonstrative Thought". *Mind and Language* 25:169–195. 2, 6, 32, 40, 44, 49, 50, 51, 52, 53, 54, 55, 56, 57, 58, 61, 66, 91, 93, 94, 98

Locke, John. 1997/1689. *An Essay Concerning Human Understanding*. Penguin Books. 119

Logothetis, Nikos K. 1998. "Single Units and Conscious Vision". *Philosophical Transactions: Biological Sciences* 353:1801–1818. 71, 105

Logue, Heather. 2009. *Perceptual Experience: Relations and Representations*. Ph.D. thesis, MIT. 26, 115, 117

——. 2012a. "What Should the Naïve Realist Say about Total Hallucinations?" *Philosophical Perspectives* 26:173–199. 25, 114, 122

——. 2012b. "Why Naïve Realism?" *Proceedings of the Aristotelian Society* 112:211–237. 25, 46, 74, 114, 118, 126

——. 2014. "Experiential Content and Naive Realism: A Reconciliation". In Berit Brogaard (ed.), *Does Perception Have Content?* Oxford University Press, 220–41. 120

Lowe, E.J. 2007. "Sortals and the Individuation of Objects". *Mind and Language*

22:514–33. 40

Lycan, William G. 2001. "The Case for Phenomenal Externalism". *Philosophical Perspectives* 15:17–35. 21

Machery, Edouard. 2009. *Doing without Concepts*. Oxford University Press. 7, 33, 55

Mack, Arien and Rock, Irv. 1998. *Inattentional Blindness*. The MIT Press. 16, 108

Marr, David. 1980. "Visual Information Processing: The Structure and Creation of Visual Representations". *Phil Trans R. Soc London B* 290:199–218. 76, 77, 78, 84

——. 1982. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. New York: Freeman. 12, 76

Marr, David and Nishihara, Herbert K. 1978. "Representation and recognition of the spatial organization of three-dimensional shapes". *Proc. R. Soc. Lond. B* 200:269–294. 77

Martin, M. G. F. 2002a. "Particular Thoughts and Singular Thought". In Anthony O'Hear (ed.), *Logic, Thought, and Language*, volume 51 of *Royal Institute of Philosophy Supplements*. Cambridge University Press, 173–214. 1, 7, 14, 21, 22, 27, 30, 31, 32, 34, 35, 37, 49, 53, 54, 93, 118, 121

——. 2002b. "The Transparency of Experience". *Mind and Language* 4:376–425. 118

——. 2004. "The Limits of Self-Awareness". *Philosophical Studies* 120:37–89. 37, 46, 74, 116

——. 2006. "On Being Alienated". In Tamar Szabo Gendler and John Hawthorne (eds.), *Perceptual Experience*. Oxford University Press, 354–410. 118

McDowell, John. 1982. "Truth-Value Gaps". In L. Jonathan Cohen, Jerzy Loś, Helmut Pfeiffer, and Klaus-Peter Podewski (eds.), *Logic, Methodology and Philosophy of Science VI*, volume 104 of *Studies in Logic and the Foundations of Mathematics*. Elsevier, 299–313. 27, 121

——. 1984. "De Re Senses". *Philosophical Quarterly* 34:283–294. 27, 121

——. 1986. "Singular Thought and the Extent of Innter Space". In John McDowell and Philip Pettit (eds.), *Subject, Thought, and Context*. Oxford University Press,

137–168. 1, 3, 22, 30, 34, 35

——. 1994a. "The Content of Perceptual Experience". *Philosopical Quarterly* 44:190–205. 74

——. 1994b. *Mind and World*. Harvard University Press. 14

——. 2010. "Tyler Burge on Disjunctivism". *Philosophical Explorations* 13:243–255. 30

——. 2013. "Perceptual Experience: Both Relational and Contentful". *European Journal of Philosophy* 21:144–157. 30

McGinn, Colin. 1989. *Mental Content*. Blackwell. 14

Mehta, Neil. 2013. "Beyond Transparency: The Spatial Argument for Experiential Externalism". *Philosophers' Imprint* 13:–. 21

Mele, Alfred R. 2009. "Intention and Intentional Action". In Brian P. McLaughlin, Ansgar Beckermann, and Sven Walter (eds.), *The Oxford Handbook of Philosophy of Mind*. Oxford University Press, 691–710. 90

Millar, Boyd. 2014a. "The Phenomenological Directness of Perceptual Experience". *Philosophical Studies* 170:235–53. doi:10.1007/s11098-013-0210-y. 116, 118

——. 2014b. "The Phenomenological Problem of Perception". *Philosophy and Phenomenological Research* 88:625–54. 116, 118

Millikan, Ruth G. 1984. *Language, Thought and Other Biological Categories*. The MIT Press. 3, 50

——. 1989. "Biosemantics". *Journal of Philosophy* 86:281–97. 3

——. 2004. *Varieties of Meaning: The 2002 Jean Nicod Lectures*. MIT Press. 3, 50, 55, 123

——. 2009. "Biosemantics". In Brian P. McLaughlin, Ansgar Beckermann, and Sven Walter (eds.), *The Oxford Handbook of Philosophy of Mind*. Oxford University Press, 394–406. 3

——. 2012. "Are There Mental Indexicals and Demonstratives?" *Philosophical Perspectives* 26:217–234. 32, 93

Milner, A. David and Goodale, Melvyn A. 2006. *The Visual Brain in Action*. Oxford Psychology Series. Oxford University Press, 2 edition. 94, 104, 105

Milner, A.D. and Goodale, M.A. 2008. "Two Visual Systems Re-viewed". *Neuropsychologia* 46:774–785. 104

Mitroff, Stephen R., Scholl, Brian J., and Wynn, Karen. 2005. "The Relationship Between Object Files and Conscious Perception". *Cognition* 96:67–92. 74

Mole, Christopher, Smithies, Declan, and Wu, Wayne (eds.). 2011a. *Attention: Philosophical & Psychological Essays*. Oxford University Press. 57

Mole, Christopher, Smithies, Declan, and Wu, Wayne. 2011b. "Introduction". In Christopher Mole, Declan Smithies, and Wayne Wu (eds.), *Attention: Philosophical & Psychological Essays*. Oxford University Press. 58, 87

Moore, Cathleen M., Yantis, Steven, and Vaughan, Barry. 1998. "Object-based visual selection: evidence from perceptual completion". *Psychological Science* 9:104–110. 87

Nagel, Thomas. 1974. "What is it like to be a Bat?" *Philosophical Review* 83:435–450. Page citations to the reprint in Chalmers, *Philosophy of Mind: Classical and Contemporary Readings*. 2002. Oxford University Press. 12, 73, 94

Nanay, Bence. 2014a. "Empirical Problems with Anti-Representationalism". In Berit Brogaard (ed.), *Does Perception have Content?* Oxford University Press, 39–50. 121

——. 2014b. "The Representationalism versus Relationalism Debate: Explanatory Contextualism about Perception". *European Journal of Philosophy* –. doi: 10.1111/ejop.12085. 115

——. forthcoming. "Perceptual Representation/Perceptual Content". In Mohan Matthen (ed.), *Oxford Handbook for the Philosophy of Perception*. Oxford University Press. 14, 121

Noë, Alva. 2004. *Action in Perception*. The MIT Press. 12, 72, 74

Noë, Alva and Thompson, Evan. 2004. "Are there Neural Correlates of Consciousness?" *Journal of Consciousness Studies* 11:3–28. 74

Orlandi, Nicoletta. 2011a. "Embedded Seeing: Vision in the Natural World". *Noûs*

Published first online, December 29. 12, 72

——. 2011b. "The Innocent Eye: Seeing-as without Concepts". *American Philosophical Quarterly* 48:17–31. 12, 72

——. 2014. *The Innocent Eye: Why Vision Is Not a Cognitive Process*. Oxford University Press. 12, 72, 76

Palmer, Stephen E. 1992. "Common region: a new principle of perceptual grouping". *Cognitive Psychology* 24:436–47. 87

Palmer, Stephen E. and Beck, Diane M. 2007. "The repetition discrimination task: An objective method for studying perceptual grouping". *Perception and Psychophysics* 69:68–78. 87

Palmer, Stephen E. and Rock, Irvin. 1994. "Rethinking perceptual organization: The role of uniform connectedness". *Psychonomic bulletin & review* 1:29–55. 87

Pautz, Adam. 2009. "What are the Contents of Experiences?" *The Philosophical Quarterly* 59:483–507. 14, 15, 22, 23, 25, 27, 34, 36, 37, 47, 74, 114, 116, 121, 122, 125

——. 2010. "Why Explain Visual Experiences in Terms of Content". In Bence Nanay (ed.), *Perceiving the World*. Oxford University Press, 254–309. 3, 22, 23, 25, 35, 36, 45, 114, 115

——. 2011. "Can Disjunctivists Explain Our Access to the Sensible World?" *Philosophical Issues* 21:384–433. 120, 121, 125

——. 2013. "Is Phenomenology the Ground of Intentionality?" In Uriah Kriegel (ed.), *The Phenomenal Intentionality Research Program*. Oxford. 120, 125

Peacocke, Christopher. 1983. *Sense and Content: Experience, Thought, and Their Relations*. Clarendon Press. 6, 26

——. 1992a. "Scenarios, Concepts, and Perception". In Tim Crane (ed.), *The Contents of Experience: Essays on Perception*. Cambridge University Press, 105–135. 117

——. 1992b. *A Study of Concepts*. The MIT Press. 7, 14, 33, 55

——. 2009. "Concepts and Possession Conditions". In Brian P. McLaughlin, Ansgar Beckermann, and Sven Walter (eds.), *The Oxford Handbook of Philoso-*

phy of Mind. Oxford University Press, 437–456. 7, 33, 54

Pettit, Philip and McDowell, John (eds.). 1986. *Subject, Thought and Context*. Clarendon Press. 3

Poljac, Ervin, de Wit, Lee, and Wagemans, Johan. 2012. "Perceptual Wholes can Reduce the Consciouss Accessibility of their Parts". *Cognition* 123:308–312. 82

Posner, Michael I. 1980. "Orienting of Attention". *Quarterly Journal of Experimental Psychology* 32:3–25. 60, 63, 79

Posner, Michael I., Snyder, Charles R. R., and Davidson, Brian J. 1980. "Attention and the detection of signals". *Journal of Experimental Psychology: General* 109:160–174. 79

Price, H. H. 1932. *Perception*. Methuen. 26, 119

Prinz, Jesse. 2000. "A Neurofunctional Theory of Visual Consciousness". *Consciousness and Cognition* 9:243–259. 71, 77, 107

———. 2002. *Furnishing the Mind: Concepts and Their Perceptual Basis*. The MIT Press. 7, 33, 55

———. 2006. "Beyond Appearances: The Content of Sensation and Perception". In Tamar Szabo Gendler and John Hawthorne (eds.), *Perceptual Experience*. Oxford University Press, 434–460. 71, 107

———. 2011. "Is Attention Necessary and Sufficient for Consciousness?" In Christopher Mole, Declan Smithies, and Wayne Wu (eds.), *Attention: Philosophical & Psychological Essays*. Oxford University Press, 174–203. 45, 71, 91, 92, 94, 96, 107, 108

Prosser, Simon. 2011. "Affordances and Phenomenal Character in Spatial Perception". *The Philosophical Review* 120:475–513. 116

Putnam, Hilary. 1975. "The Meaning of 'Meaning'". In Keith Gunderson (ed.), *Language, Mind, and Knolwedge*, volume 7 of *Minnesota Studies in the Philosophy of Science*. University of Minnesota Press, 131–193. 2, 50

Pylyshyn, Zenon W. 1986. *Computation and Cognition: Towards a Foundation for Cognitive Science*. The MIT Press. 12

———. 2001. "Visual Indexes, Preconceptual Objects, and Situated Vision". *Cognition* 80:127–158. 31, 49, 66, 80, 93

——. 2003. *Seeing and Visualizing: It's Not What You Think*. The MIT Press. 41

——. 2007. *Things and Places: How the Mind Connects with the World*. The MIT Press. 2, 11, 12, 31, 40, 44, 49, 50, 51, 52, 53, 54, 55, 57, 66, 68, 74, 75, 79, 80, 91, 93

Pylyshyn, Zenon W. and Storm, R.W. 1988. "Tracking Multiple Independent Targets: Evidence for a Parallel Tracking Mechanism". *Spatial Vision* 3:179–197. 79

Raftopoulos, Athanassius and Müller, Vincent. 2006. "Nonconceptual Demonstrative Reference". *Philosophy and Phenomenological Research* 72:251–285. 2, 11, 44, 49, 50, 53, 54, 55, 56, 61, 67, 68, 93, 94, 98

Raleigh, Thomas. 2009. "Understanding How Experience "Seems"". *European Journal of Analytic Philosophy* 5:67–78. 116, 118

Recanati, François. 2010. "Singular Thought: In Defense of Acquaintance". In Robin Jeshion (ed.), *New Essays on Singular Thought*. Oxford University Press, 141–189. 11, 31, 49, 55, 56, 68, 93

——. 2012. *Mental Files*. Oxford University Press. 2, 6, 7, 11, 32, 44, 49, 52, 53, 54, 55, 56, 68, 93

Rensink, Ronald A., O'Regan, J. Kevin, and Clark, James J. 1997. "To See or Not to See: The Need for Attention to Perceive Changes in Scenes". *Psychological Science* 8:368–373. 16, 108

——. 2000. "On the Failure to Detect Changes in Scenes Across Brief Interruptions". *Visual Cognition* 7:127–145. 16, 108

Robinson, Howard. 1994. *Perception*. Routledge. 115

Roessler, Johannes. 2011. "Introduction: Perception, Causation, and Objectivity". In Johannes Roessler, Hemdat Lerman, and Naomi Eilan (eds.), *Perception, Causation, and Objectivity*. Oxford University Press, 1–18. 3, 19

Ronconi, Luca, Basso, Demis, Gori, Simone, , and Facoetti, Andrea. 2014. "TMS on Right Frontal Eye Fields Induces an Inflexible Focus of Attention". *Cerebral Cortex* 24:396–402. 104

Ruff, Christian C. 2011. "A Systems-Neuroscience View of Attention". In Christopher Mole, Declan Smithies, and Wayne Wu (eds.), *Attention: Philosophical &*

Psychological Essays. Oxford University Press, 1–23. 58

Russell, Bertrand. 1956. "On the Nature of Acquaintance". In Robert C. Marsh (ed.), *Logic and Knowledge*. George Allen & Unwin, 125–174. 3

———. 1997/1912. *The Problems of Philosophy*. Oxford University Press. With a New Introduction by John Perry. 26, 56, 119

Schall, Jeffrey D. 1995. "Neural Basis of Saccade Target Selection". *Reviews in the Neurosciences* 6:63–85. 106

———. 2004. "On the role of frontal eye field in guiding attention and saccades". *Vision Research* 44:1453–1467. 104

Schall, Jeffrey D., Morel, Anne, King, Dana J., and Bullier, Jean. 1995. "Topography of Visual Cortex Connection with Frontal Eye Field in Macaque: Convergence and Segregation of Processing Streams". *The Journal of Neuroscience* 15:4464–4487. 105, 106

Schellenberg, Susanna. 2010. "The Particularity and Phenomenology of Perceptual Experience". *Philosophical Studies* 149:19–48. 25, 114

———. 2011. "Perceptual Content Defended". *Noûs* 45:714–750. 21, 23, 25, 27, 29, 36, 47, 74, 114, 116, 121, 125

———. 2013. "Perceptual Content and Relations". *Philosophical Studies* 163:49–55. 25, 114, 122

Scholl, Brian J. 2001. "Objects and Attention: the State of the Art". *Cognition* 80:1–46. 41, 44, 63, 77, 78, 79, 80, 83, 84, 87, 88, 91

———. 2009. "What Have We Learned about Attention from Multiple-Object Tracking (and Vice Versa)?" In Don Dedrick and Lana Trick (eds.), *Computation, Cognition, and Pylyshyn*. The MIT Press, 49–78. 41, 66, 80

Scholl, Brian J. and Pylyshyn, Zenon W. 1999. "Tracking multiple items through occlusion: Clues to Visual Objecthood". *Cognitive Psychology* 38:259–290. 87

Scholl, Brian J., Pylyshyn, Zenon W., and Feldman, Jacob. 2001. "What is a Visual Object? Evidence from Target Merging in Multiple Object Tracking". *Cognition* 80:159–177. 41, 72, 76, 77, 78, 79, 80, 81, 82

Searle, John R. 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge University Press. 47, 53

Sekuler, Allison B. and Bennett, Patrick J. 2001. "Generalized common fate: grouping by common luminance changes". *Psychological Science* 12:437–444. 87

Sheinberg, D. L. and Logothetis, Nikos K. 1997. "The Role of Temporal Cortical Areas in Perceptual Organization". *Proceedings of the National Academy of Sciences* 94:3408–3413. 105

Siegel, Susanna. 2006a. "How Does Visual Phenomenology Constrain Object-Seeing?" *Australasian Journal of Philosophy* 84:429–441. 31, 32, 49, 53, 95

——. 2006b. "Subject and Object in the Contents of Visual Experience". *Philosophical Review* 115:355–88. 26, 115

——. 2007. "How Can We Discover the Contents of Experience?" *Southern Journal Of Philosophy* 45:127–42. 100

——. 2010a. *The Contents of Visual Experience*. Philosophy of Mind Series. Oxford University Press. 14, 21, 23, 25, 26, 29, 31, 36, 37, 49, 74, 114, 115, 116, 121, 123, 125

——. 2010b. "Do Visual Experiences Have Contents". In Bence Nanay (ed.), *Perceiving the World*. Oxford University Press, 333–368. 25, 114, 117, 120, 121, 125

Siegel, Susanna and Silins, Nicholas. forthcoming. "The Epistemology of Perception". In Mohan Matthen (ed.), *Oxford Handbook of Philosophy of Perception*. Oxford University Press, –. 14

Siewert, Charles. 1998. *The Significance of Consciousness*. Princeton University Press. 120, 125

——. 2012. "Respecting Appearances: a Phenomenological Approach to Consciousness". In Dan Zahavi (ed.), *The Oxford Handbook of Contemporary Phenomenology*. Oxford University Press. 12

Singh, Manish and Hoffman, Donald D. 2001. "Part-Based Representations of Visual Shape and Implications for Visual Cognition". In Thomas F. Shipley and Philip J. Kellman (eds.), *From Fragments to Objects: Segmentation and Grouping in Vision*, volume 130 of *Advances in Psychology*. North-Holland, 401–459. 84

Singh, Manish and Scholl, Brian J. 2000. "Using attentional cueing to explore

part structure". Poster presented at the 2000 Pre-Psychonomies Object Perception and Memory meeting, New Orleans, LA. 84

Smith, A.David. 2002. *The Problem of Perception*. Harvard University Press. 115

Smithies, Declan. 2011a. "Attention is Rational-Access Consciousness". In Christopher Mole, Declan Smithies, and Wayne Wu (eds.), *Attention: Philosophical & Psychological Essays*. Oxford University Press, 247–273. 32, 49, 61, 95

——. 2011b. "What is the Role of Consciousness in Demonstrative Thought?" *Journal of Philosophy* 108:5–34. 5, 17, 18, 19, 31, 32, 34, 49, 51, 53, 93, 95, 98, 110, 112

Sollberger, Michael. 2012. "Causation in Perception: A Challenge to Naïve Realism". *Review of Philosophy and Psychology* 3:581–95. 118

Soteriou, Matthew. 2000. "The Particularity of Visual Perception". *European Journal of Philosophy* 8:173–189. 27, 47

Speaks, Jeff. 2009. "Transparency, Intentionalism, and the Nature of Perceptual Content". *Philosophy and Phenomenological Research* 79:539–573. 23, 25, 27, 31, 36, 47, 49, 53, 74, 114, 122

Spelke, Elizabeth S. 1990. "Principles of Object Perception". *Cognitive Science* 14:25–56. 84

Sperling, George. 1960. "The Information Available in Brief Visual Presentations". *Psychological Monographs: General and Applied* 74:1–29. 108

Stanton, Gregory B., yi Deng, Shu, Goldberg, Michael E., and McMullen, Nathaniel T. 1989. "Cytoarchitectural Characteristic of the Frontal Eye Fields in Macaque Monkeys". *The Journal of Comparative Neurology* 282:415–427. 103

Taylor, Kenneth. 2010. "On Singularity". In Robin Jeshion (ed.), *New Essays on Singular Thought*. Oxford University Press, 77–102. 55

Thompson, Kirk G. and Bichot, Narcisse P. 2005. "A visual salience map in the primate frontal eye field". *Progress in Brain Research* 147:251–262. 104

Thompson, Kirk G., Biscoe, Keri L., and Sato, Takashi R. 2005. "Neuronal Basis of Covert Spatial Attention in the Frontal Eye Field". *The Journal of Neuroscience* 25:9479–9487. 104

Travis, Charles S. 2004. "The Silence of the Senses". *Mind* 113:57–94. 14, 26, 121, 125

Treisman, Anne. 1998. "Feature Binding, Attention and Object Perception". *Phil Trans R. Soc London B* 353:1295–1306. 20, 77

Treisman, Anne and Gelade, Garry. 1980. "A Feature-Integration Theory of Attention". *Cognitive Psychology* 12:97–136. 20, 77

Tye, Michael. 1995. *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. The MIT Press. 3, 15, 25, 27, 47, 71, 74, 94, 114, 116, 122

Ungerleider, Leslie G. and Mishkin, Mortimer. 1982. "Two Cortical Visual Systems". In David J. Ingle, Melvyn A. Goodale, and Richard J.W. Mansfield (eds.), *Analysis of Visual Behavior*. The MIT Press, 549–586. 104

Valberg, Jerome J. 1992a. *The Puzzle of Experience*. Oxford University Press. 115, 116

Valberg, J.J. 1992b. "The Puzzle of Experience". In Tim Crane (ed.), *The Contents of Experience: Essays on Perception*. Cambridge University Press, 18–47. 30, 32, 33, 44, 49, 91, 95

van Gelder, Tim. 1995. "What might cognition be if not computation?" *Journal of Philosophy* 92:345–81. 12, 72

Wagemans, Johan, Elder, James H., Kubovy, Michael, Palmer, Stephen E., Peterson, Mary A., Singh, Mansih, and von der Heydt, Rüdiger. 2012. "A Century of Gestalt Psychology in Visual Perception: I. Perceptual Grouping and Figure-Ground Organization". *Psychological Bulletin* 138:1172–1217. 71, 77, 78, 83, 84, 85, 86, 87

Walker, A. Earl. 1940. "A Cytoarchitectural Study of the Prefrontal Area of the Macaque Monkey". *Journal of Comparative Neurology* 73:59–86. 103

Wandell, Brian A. and Winawer, Jonathan. 2011. "Imaging retinotopic maps in the human brain". *Vision Research* 51:718–737. 63

Wang, Lan, Weng, Xuchu, and He, Sheng. 2012. "Perceptual Grouping without Awareness: Superiority of Kanizsa Triangle in Breaking Interocular Suppression". *PLoS ONE* 7:e40106. 78

Watson, Stephen E. and Kramer, Arthur F. 1999. "Object-based visual selective

attention and perceptual organization". *Perception & Psychophysics* 61:31–49.
87

Watzl, Sebastian. 2011. "Attention as Structuring of the Stream of Consciousness". In Christopher Mole, Declan Smithies, and Wayne Wu (eds.), *Attention: Philosophical & Psychological Essays*. Oxford University Press, 145–173. 57, 58, 87

Weiskrantz, Lawrence. 2009. *Blindsight: A Case Study Spanning 35 Years and New Developments*. Oxford Psychology Series. Oxford University Press. 16, 17, 97

Wertheimer, Max. 1923. "Untersuchungen zur Lehre von der Gestalt, II. [Investigations in Gestalt Theory: II. Laws of organization in perceptual forms]". *Psychologische Forschung* 4:301–350. 84, 86

Wu, Wayne. 2011a. "Attention as Selection for Action". In Christopher Mole, Declan Smithies, and Wayne Wu (eds.), *Attention: Philosophical & Psychological Essays*. Oxford University Press, 97–116. 44, 56, 57, 58, 87, 91, 94, 98

———. 2011b. "What is Conscious Attention?" *Philosophy and Phenomenological Research* 82:93–120. 31, 32, 49, 53, 56, 57, 58, 87, 94, 95, 98