

ON LEARNING FROM EXPERIENCE

by

Ori Beck

B.A. in Philosophy and Criminology, Bar-Ilan University, 2005

M.A. in Philosophy, Tel-Aviv University, 2008

M.A. in Philosophy, University of Pittsburgh, 2015

Submitted to the Graduate Faculty of
the Kenneth P. Dietrich School of Arts and Sciences in partial
fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2015

UNIVERSITY OF PITTSBURGH
KENNETH P. DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Ori Beck

It was defended on

September 9, 2015

and approved by

Anil Gupta, Alan Ross Anderson Distinguished Professor of Philosophy

Edouard Machery, Professor of History and Philosophy of Science

John McDowell, Distinguished University Professor of Philosophy

Mark Wilson, Professor of Philosophy

Dissertation Director: Anil Gupta, Alan Ross Anderson Distinguished Professor of
Philosophy

Copyright © by Ori Beck
2015

ON LEARNING FROM EXPERIENCE

Ori Beck, PhD

University of Pittsburgh, 2015

How can I come to justifiably (or knowledgeably) believe that there is an owl over there on the basis of my seeing its being there? More generally: How can we come to justifiably (or knowledgeably) believe things on the basis of our experiences? I argue that we can make do with a simple and intuitive answer: Experiences present us with objects, properties, relations and states of affairs (or, events). If we have the capacities to tell what these presented entities are, we can justifiably (or knowledgeably) form appropriate beliefs on the basis of the experiences. I call this position the “Simple Picture” (or, “SP”).

SP stands in contrast to the Representationalist epistemology, according to which experiences do not quite present us with entities, as represent how the world is. Under certain conditions, we can justifiably form beliefs on the strength of these experiences by taking it that things are as our experiences represent them to be.

The Representationalist epistemology cannot be correct. Through a series of thought experiments, involving various kinds of spectrum inversions, I show that none of the current theories of content can assign representational contents to our experiences in a way that tracks the beliefs that we can justifiably form on their strength. I then further show that even intuitive and pre-theoretical attributions of representational contents cannot perform this feat.

Representationalist epistemologists might retort by saying that SP must have problems of its own: “SP holds that perceptual experiences are mental states that present us with entities. This means that it belongs to the Naive Realist family of views. And all views of this family cannot give an adequate account of illusions and hallucinations.” SP’s response sets it apart

from other Naive Realist views. While most Naive Realist views (e.g., those of Brewer, Fish and Martin) favor a Metaphysical Disjunctivist line, SP offers a uniform metaphysics and epistemology for perceptual and hallucinatory experiences. SP furthermore gives a positive account of illusions as well.

TABLE OF CONTENTS

PREFACE	ix
1.0 A BRIEF INTRODUCTION	1
2.0 HOW <i>NOT</i> TO LEARN FROM EXPERIENCE	4
2.1 On ER	5
2.2 A Spectrum Inversion Argument	7
2.3 Experience-First Accounts	19
2.4 Function-First Accounts	25
2.5 A Last Line of Defense	31
2.6 A Look Ahead	32
3.0 HOW WE LEARN FROM PERCEPTION	34
3.1 Our Implicit Conception	35
3.2 On Perceptual Experiences	38
3.3 On Capacities to Tell	41
3.3.1 “Owing To”	47
3.3.2 “Inferentially Coming to Know”	47
3.3.3 “Experiential Tracking Effort”, “Success” and “Control”	48
3.3.4 “Successful Exercise”	54
3.4 On Knowing, Justifiably Believing, and Rationally Believing	57
3.5 A Test Case	59
3.6 Conclusion	64
4.0 OBJECTS OF HALLUCINATION	66
4.1 Troubles for Hermeneutic Irrealism	69

4.1.1	A First Charge for Realism	69
4.1.2	Three Hermeneutic Irrealisms	72
4.1.3	The Problem of Ignorance and Error	74
4.1.4	The Problem of Communication	77
4.1.5	The Problem of Theoretical and Practical Inference	79
4.1.6	A Reply?	82
4.2	A Realist Proposal	87
4.2.1	Abstract Realism	87
4.2.2	The Nonexistence Objection	94
4.3	For Abstract Realism	97
4.3.1	A Default Theory	97
4.3.2	The Phenomenology of Hallucinations	98
4.3.3	The Epistemology of Hallucinations	100
4.3.4	The Semantics of Hallucinations	106
4.4	Responses to Doubts	110
5.0	ON LEARNING FROM HALLUCINATION AND ILLUSION	119
5.1	On Hallucinations and Other Experiences	120
5.2	On Illusions	125
5.3	Conclusion	134
6.0	Bibliography	136

LIST OF FIGURES

1	Kanizsa triangle	131
---	----------------------------	-----

PREFACE

Stanley Cavell (1979, p. 125) once described philosophy thus:

In the face of the questions posed in Augustine, Luther, Rousseau, Thoreau . . . , we are children; we do not know how to go on with them, what ground we may occupy. In this light, philosophy becomes the education of grownups.

I am not sure if all my philosophy teachers would accept Cavell's description of their trade. Nor am I much of a grownup. But Cavell's description is surely right insofar as it implies that my teachers were also my educators. For this I will always be grateful to them.

Thank you to Anil Gupta. Anil taught me to explore bold and unconventional ideas, to enjoy the intricacies of difficult puzzles, and to never lose sight of all the constraints along the way. He has made an indelible mark on my thinking. The force and clarity of his ideas shaped every part of this essay, and without him it would have never been written. Not only was Anil a philosophical mentor to me, he also educated me in other ways. As a logician, Anil never failed to insist that the variables occurring in my definienda also made their way into my definienda. And as food critics, he and his wife Mukta are responsible for much of my love of Indian food (especially spicy samosas).

Thank you to Hilary Putnam. I first met Hilary more than 10 years ago, when I was still working on my B.A. It was out of pure generosity that he agreed to meet me. We have kept in regular touch since. During our long walks around Tel Aviv University and along the city's promenade, Hilary made me a realist and a semantic externalist. He taught me about logic, mathematics, language, metaphysics and even ethics. I cannot imagine philosophy without his ideas. Hilary also instructed me in English poetry, on the correct way of drinking cappuccino (no sugar!), and corrected my pronunciation of "gigabyte".

Each in his own way, both Anil and Hilary showed me how one can be tender, generous, kind, and personal in doing philosophy. The forthrightness, courage and incisiveness of Ruth Anna, Hilary's wife, are traits I will forever try to emulate.

Thank you also to John McDowell, Edouard Machery and Mark Wilson. I treasure every minute I spent with them and their writings. They have always challenged me and pushed me further. By working sometimes with and (to my great distress) sometimes against their ideas, I learnt much about how to philosophize. I admire their thinking immensely. Their books—*Mind and World*, *Doing Without Concepts*, and *Wandering Significance*—are never far from my mind.

Other faculty at the University of Pittsburgh and at Carnegie Mellon University also provided me with incredibly helpful lessons and discussions. Thank you especially to Wayne Wu, Thomas Ricketts, Mazviita Chirimuuta, James Shaw, James G. Lennox, and Kenneth F. Schaffner. Thank you also to Christopher Hill, who visited the University of Pittsburgh and delivered an astonishingly illuminating seminar on the Philosophy of Mind.

Lady Luck also contributed to my education. A lucky dinner conversation with Bill Brewer changed my mind on the Myth of the Given. Separate encounters with Alva Noë and Declan Smithies redirected my thoughts about phenomenology. John Campbell's visit to Pittsburgh set me thinking about demonstrative reference. And a breakfast conversation with Graham Priest at a Brazilian hotel made me see the appeals of Noneism. Helpful comments from Martine Nida-Rümelin and Diana Raffman (who were also there in Brazil) made their mark on chapter 2.

Of my fellow graduate students, Jonathan Buttaci, David de Bruijn, Ulf Hlobil, Adam Marushak, Joe Milburn, Eric Palmer, Shivam Patel, Raja Rosenhagen, Jonathan Surovell, Laura Tomlinson and Miloš Vuletić all deserve a special thanks. David and Ulf are especially to be thanked for resisting me at every turn.

Before I started my studies at the University of Pittsburgh, my intellectual life was guarded by many wonderful Israeli teachers. I am particularly grateful for the instruction of Yitzhak Benbaji, Marcelo Dascal, Eli Dresner, Charlotte Katzoff, Anat Matar, Yaron Senderowicz, Ofer Shamai, and Ruth Weintraub.

Anat Biletzki, my M.A. advisor at Tel Aviv (now at Quinnipiac University), taught me not only philosophy, but also politics and civil rights. She made a great impact on my development. Her views, wit and level-headed attitude fill me with joy every time we speak.

David Widerker took me under his wing when I was a young B.A. student. As my teacher, he unrelentingly drilled me into thinking carefully and consistently. As my friend, he is always there for me. A true hero.

Shlomo Giora Shoham sent me to college when I was only a Bar Mitzvah boy. Without his intervention, I would probably have become a lawyer. Instead, he educated me on Existentialism, world mythology, the Pentateuch and the creative process. There is no measure to what I owe him.

Last, but not least, I thank those closest to my heart:

Alissa, my sweetheart. You shared in every single day of thinking and writing that went into this essay. Your encouragement, love and support nourished and comforted me. Your part in my work is larger than you can imagine.

Sapir and Amnon, my beloved parents. Without your love and devotion, not. I dedicate this essay to you. May we forever be together.

1.0 A BRIEF INTRODUCTION

There is, nowadays, a duel going on between an orthodox and a heterodox view of the nature of perceptual experience. The orthodox view is called “Representationalism” (or “Intentionalism”). It says that perceptual experiences are *representational* mental states, i.e., mental states evaluable for truth (veridicality, accuracy) at (centered) possible worlds.¹ The heterodox view is called “Relationalism” (or, “Naive Realism”). This view denies that perceptual experiences are representational states. Instead, it says that they are *presentational* mental states, i.e., mental states that immediately present worldly items to whoever has them.² And while proponents of either view argue that they can offer the best accounts of the metaphysics, phenomenal character, epistemic powers and biological underpinnings of perceptual experience, no consensus seems to be emerging.

Perhaps the lack of consensus is a good thing. For there is always the possibility that both views have something to them. In other words, perhaps perceptions *both* immediately present worldly items to the subject, *and* are evaluable for truth (veridicality, accuracy) at (centered) possible worlds. Absent an argument to rule out this possibility, it seems premature to try to settle the duel between representationalists and relationalists. And as I

¹Representationalism is championed by Tyler Burge (1986; 2003a; 2010), Martin Davies (1992; 1997), Gareth Evans (1982), Gilbert Harman (1990), Christopher Hill (2009), John McDowell (1996), Colin McGinn (1991), Ruth Garrett Millikan (2004), Adam Pautz (2010), Christopher Peacocke (1992; 2001; 2004), James Pryor (2000), John Searle (1983), Susanna Siegel (2010) and Michael Tye (1995; 2000; 2009). In his (2009), John McDowell updated his view of perceptual experience from (1996) to suggest that the content of such experiences can be given by complex demonstrative phrases of the form *THAT* – *F* (where the concepts occurring in *F* may include the proper and common sensibles). Though this view is not strictly speaking a representational one, it is in all relevant ways sufficiently close to representationalism to count as a representationalist view for the purposes of this essay.

²Relationalism is championed by William Alston (1999), Bill Brewer (2011), John Campbell (2002; 2009; 2014), Matthew Conduct (2012), William Fish (2009), James Genone (2014), Anil Gupta (2012), John M. Hinton (1967), Mark Johnston (2004; 2006), Michael G. F. Martin (2002; 2004), David (A. D.) Smith (2002), Paul Snowdon (1981; 1990) and Charles Travis (2013a).

know of no such argument, I will not try to settle it.

What I try to settle instead is a related matter: How can we come to justifiably (or knowledgeably) believe things on the basis of our experiences? Should we prefer a representationalist or a relationalist account? I argue that, when it comes to this specific matter, the best and most intuitive account is wholly relationalist. In a nutshell, this account (call it “the Simple Picture”, or “SP”, for short) says that experiences present us with objects, properties, relations and states of affairs (or, events). If we have the capacities to tell what these presented entities are, we can justifiably (or knowledgeably) form appropriate beliefs on the basis of the experiences. This entirely intuitive account is relationalist, since it relies on the presentational nature of experience. The account is not representationalist, since it is entirely silent on whether experiences are representational states.

Representationalists will, of course, resist. They will remind me that they have their own account of how we learn from experience. They might even suggest that my relationalist account can be reduced to (or subsumed under) theirs.

In response, I show that the representationalist account—which I call “Epistemic Representationalism” (or, “ER”)—cannot be correct. ER suggests that the truth (veridicality, accuracy) conditions of experiences accord, in a certain sense, with those of the beliefs we can justifiably form on experiential grounds. Through a series of thought experiments, involving various kinds of spectrum inversions, I show that no plausible theory for assigning truth (veridicality, accuracy) conditions to experiences can accommodate this accordance. I then further show that even intuitive and pre-theoretical assignments of truth (veridicality, accuracy) conditions to experiences cannot accommodate it. I conclude from this that ER should be abandoned.

A different complaint that representationalists will make about SP is that it, like most relationalist accounts, breaks down when it comes to illusions and hallucinations. In particular, they will complain that a hallucinator can have a hallucination in which she is presented with no entities at all. In such a case, the hallucinator’s capacities to tell what entities she is presented with will have to remain idle. But by SP’s lights, if the capacities remain idle, the hallucinator will not be able to justifiably form any beliefs on the strength of her hallucination. Since the hallucinator can form such beliefs, SP must be false or incomplete.

My response to this and to related objections sets SP apart from most other relationalist views. Most relationalists accept metaphysical disjunctivism - a view on which perceptions do, but hallucinations do not, present entities to the subject.³ As a result, such relationalists must accept the first premise of the representationalist argument from the previous paragraph. In contrast with them, I hold that hallucinations and perceptions are alike in that both present entities to the subject. By thus rejecting metaphysical disjunctivism, I can use SP to provide a unified account of how we come to justifiably (or knowledgeably) believe things on the basis of both perceptual and hallucinatory experiences.

To be more specific, my view is that if you are hallucinating your mother, you are thereby presented with a perfectly ordinary object. If you are hallucinating a giant spaghetti monster, you are thereby presented with a mind-dependent hallucinated object. You can refer to both in thinking “that is my mother I am hallucinating”, or “that must be a figment of my imagination”. In articulating this view (I call it “Abstract Realism”), I explain the semantics of thoughts and assertions concerning hallucinations. I further explain that accepting an ontology of mind-dependent hallucinated entities is no more unnatural than accepting an ontology of other mind-dependent entities, such as nations or monetary values.

With respect to illusions, I claim (i) that illusions are a kind of seeming-state (i.e., a mental state that consists of its seeming to one that something is the case), and (ii) that seeming-states are not experiences, but something that experiences (whether perceptual or hallucinatory) lead to. Since illusions are not experiences, the question of the epistemic powers of illusions (if such there be) falls beyond the scope of SP and this essay.

The structure of this essay is as follows: After this introductory chapter, chapter 2 argues that ER must be rejected. Chapter 3 then presents the parts of SP which concern perception based beliefs. Chapter 4 then leaves the epistemological thread of the discussion, and argues for my view of hallucinations - abstract realism. Chapter 5 then concludes the essay, by presenting the parts of SP which concern hallucination based beliefs, and by discussing the topic of illusions. Chapters are written so as to be readable in any order.

³ Champions of metaphysical disjunctivism are Bill Brewer (2011), John Campbell (in his and Quassim Cassam’s 2014), William Fish (2009), and Michael G. F. Martin (2002; 2004).

2.0 HOW *NOT* TO LEARN FROM EXPERIENCE

Suppose I see a red couch at the side of the room. Seeing it, I can come to justifiably believe that it is red. But how does my sighting support my belief? More generally: How can we come to justifiably believe this or that on the strength of experiences? And why is it that only certain beliefs can be justifiably had on the strength of a given experience? These questions are arguably the most fundamental questions at the intersection of epistemology and the philosophy of mind. They deserve informative answers.

Contemporary philosophy has an orthodox answer to these questions. Martin Davies (1992, p. 22-23) puts it as follows: “An experience may present the world to a subject as containing something square in front of her; and the subject may take that experience at face value and judge that there is something square in front of her.” John McDowell (1996, p. 26) agrees: “That things are thus and so is the content of the experience, and it can also be the content of a judgment: it becomes the content of a judgment if the subject decides to take the experience at face value.” So despite their many differences, McDowell and Davies both accept the orthodox view of how we learn from experience, which relies on the idea that we can “take” experiences “at their face values”. When this view is distilled from the writings of the various thinkers who further articulate it—thinkers as diverse as Tyler Burge (2003a), John McDowell (1996), Christopher Peacocke (1992; 2004) and James Pryor (2000)—we get what I will call “Epistemic Representationalism” (or, “ER” for short).

This chapter will argue that ER must be abandoned. I begin by explaining, in the section 2.1, what ER says. Sections 2.2-2.4 then deploy three thought experiments which together show that on theoretically plausible ways of assigning contents to experiences, ER turns out to be false. The first thought experiment is a version of the familiar inverted spectrum scenario (presently used to combat not a metaphysical position, but an epistemological or-

thodoxy). The second thought experiment is a more sophisticated version of the first, and the third thought experiment is what we might call a “functional inversion” scenario. After these sections, section 2.5 argues against an attempt to defend ER on intuitive grounds. The final section 2.6 briefly sketches the view that I believe should take ER’s place in the epistemological landscape.

2.1 ON ER

To see what ER says, begin with the thought that ER accepts Representationalism - the view that any experience has a *content*, or a *face value* - a way that things are according to the experience.¹ These face values make experiences evaluable for truth (veridicality, accuracy) at (centered) possible worlds.² [For example, if I see a red couch, the face value of my experience might be “that is a red couch”. This experience is veridical at a world w iff what I actually see is indeed a red couch at w .³] And since all experiences are evaluable in such ways, we may match any experience with an intension, i.e., a function assigning the experience a veridicality value at any (centered) possible world.⁴

That experiences can be matched with intensions fits nicely with another aspect of ER’s. According to ER, one can form beliefs by *taking* our experiences *at* their face values. And, if you come to believe something by taking an experience at its face value, then it is impossible

¹The name “Representationalism” (as well as the name “Intentionalism”) is sometimes reserved for the view that (i) experiences have contents, and (ii) the phenomenal character of an experience supervenes on its content. To avoid confusions, let us call this view “phenomenal Representationalism”. Clearly, one can be a representationalist without being a phenomenal representationalist..

²Henceforth, when I remark that an experience is veridical (falsidical), the reader should take that to mean that the experience is veridical, true, or accurate (falsidical, false, or inaccurate).

³Versions of ER that evaluate experiences at centered possible worlds might alternatively suggest that my experience is veridical at $\langle w, i \rangle$ iff what is appropriately related to i at w is a red couch at w .

⁴Two remarks: (1) Note that we may match experiences with intentions without taking any stand about whether the face values of the experiences are structured or unstructured entities, or about whether the face values engage one’s concepts or conceptual capacities. (2) John McDowell has, in his (2009), updated his view of experience from (1996). His updated position is that the content of experiences can be given by complex demonstrative phrases of the form *THAT* – *F* (where the concepts occurring in *F* may include the proper and common sensibles). So read, the view is not strictly speaking a representational one, because on it, experiences are not assigned what we have defined as intensions, but rather what we might call “subintensions” - (possibly partial) functions that map any possible world w to the experience’s referent at w . This fact might make it difficult to see how McDowell’s updated view constitutes a version of ER. This hurdle, however can be overcome. I will not trouble over this here.

that things are the way your experience portrays them to be but the proposition you thereby come to believe is false. [Example: My experience portrays something to be a red couch. By taking my experience at its face value I can come to believe that thing to be red. And ER informs me that I can only form this belief in this way given that the thing cannot both be a red couch (as the experience has it) and fail to be red (as the belief has it).] A slightly pithier way of putting ER’s commitment is the following. Define “experience e entails that p ” thus: Experience e entails that p iff there is no possible world w such that e is veridical at w , but the proposition that p is false at w .⁵ ER’s claim now becomes the claim that if you come to believe that p by taking an experience e at its face value, then e entails that p .

We now come to ER’s first core commitment. According to ER, there is a familiarly *direct* way of coming to believe things on the strength of experiences. Whenever we come to justifiably believe things in this direct experience-based way, we end up taking our experiences at their face values. We therefore have:

(ER1) *If one directly comes to justifiably believe that p on the strength of an experience e , then e entails that p .*

Subscribers to ER elucidate what they mean by a “direct” way of forming beliefs thus: They say that a person forms a belief *indirectly* on the strength of an experience iff the person forms it on the strength of both the experience and some extra-experiential resources, such as collateral beliefs. A person forms a belief *directly* on the strength of an experience iff no such extra-experiential resources are involved.⁶

The second of ER’s core commitments is to the idea that if you directly come to believe that p on the strength of an experience e , then (barring irrationality or the presence of

⁵This definition will of course require some special tailoring to fit views on which the intension of an experience (or a proposition) is a function defined over centered worlds, a partial function, or a function that can assign more than two possible veridicality values. Still, ER’s crucial point remains: There can be semantic relations of veridicality preservation between the intension of an experience and the intension of a belief.

⁶Note that many supporters of ER—e.g., [McDowell, 2009](#); [Peacocke, 2004](#); [Pryor, 2000](#)—distinguish between *inferentially* and *indirectly* forming a belief on the strength of an experience. They are moved to make this distinction by considering cases of the following kind: Say Alice is an expert at recognizing Hondas by sight, and that one day she exercises her expertise in coming to know, of some car, that it is a Honda. Since Alice is an expert at visually recognizing Hondas, some supporters of ER find it implausible to maintain that Alice came to have her knowledge inferentially. On the other hand, they also hold that Alice’s experiences cannot represent something as a Honda and, therefore, that the grounds for Alice’s knowledge extend beyond her experience. Thus, Alice’s knowledgeable belief turns out to be formed both non-inferentially and indirectly.

defeaters) your belief would be justifiably had if e entails that p . In other words, we have:

(ER2) *If (i) e entails that p , and (ii) one directly comes to believe that p on e 's strength (where one is rational and has no defeaters), then one justifiably believes that p .*

I have no objection to the idea that experiences have contents, or face values. But I do *not* think that these contents can perform the epistemological services the orthodoxy requires of them. So I believe that ER must be abandoned. I will now begin to explain why.

2.2 A SPECTRUM INVERSION ARGUMENT

Consider the scenario *Alice*, a variant on the intrapersonal spectrum inversion scenario:⁷ At stage 1 Alice—someone with normal color vision—is 18 years old and broke. In need of funds, she agrees to undergo a red-green “Experiential Color spectrum Inversion” operation (an “ECI-op”, for short). During this operation, a chip is inserted into her optic nerve, where it replaces signals that pre-op were typically caused by red and reddish (green and greenish) things with those that pre-op were typically caused by green and greenish (red and reddish) things. The effect of the operation is this: Post-op, red and reddish things (in any lighting conditions C) look to Alice the colors that green and greenish things (in C) looked to her pre-op. Similarly, green and greenish things (in C) look to her the colors that red and reddish things (in C) looked to her pre-op.⁸ Before surgery, the doctors teach Alice about the ECI-op’s expected effect in great detail. In fact, they teach her, for any shade S (with its characteristic hue, saturation and lightness) that an object can have, and for any lighting conditions C , what shade things of shade S (in C) will look to her post-op.⁹

⁷I am here heavily indebted to presentations of the interpersonal spectrum inversion scenario given by Ned Block (1990; 2007c) and Sydney Shoemaker (1982; 2007).

⁸This point is a good opportunity to mention that in this chapter I am presupposing that objects can truly be said to have colors (color properties). The presupposition, however, is not essential for my argument. Readers who cannot stomach it should feel free to replace my talk about things being colored with talk about things being spatially located, and my talk of a red-green spectrum inversions with talk of an upside-down location inversions.

⁹Imagine the doctors accomplish this using a large collection of Munsell color chips, matching chips that represent colored objects with chips that represent how those colors would look to her post-op in various lighting conditions.

At stage 2 Alice wakes up from her ECI-op. Things look strangely colored to her, and she reports things like “the fire truck outside looks to me now the color fresh grass used to look to me”. Still, since the doctors taught her, for any S and C , how things of shade S (in C) will look to her post-op, she can tell what colors things are by sight. To illustrate how she performs this feat, consider a case in which she is faced with a red couch under good light. She reasons as follows: “This couch looks to me in this light the color green things looked to me in this light pre-op; and I know that in this light red things look to me the color green things used to look to me in this light pre-op. So the couch is red.” Similar forms of reasoning allow her to accurately tell by sight not only the hue of things, but also their color saturation, their lightness, and their level of warmth or cool. She can even tell which objects are similar and dissimilar in color. For example she can reason as follows: “Objects X, Y and Z look to me in this light the color green, red and cyan things (respectively) used to look to me in this light pre-op. I know that in this light red, green and purple things look to me the color green, red and cyan things (respectively) used to look to me in this light pre-op. So X, Y and Z must be red, green and purple (respectively). So Z is more similar in color to X than to Y.” Succinctly put: Post-op, Alice’s capacity to make color judgments—i.e., make judgments about the hues, saturation levels, lightness levels, warmth levels and color similarities of things—will be nearly indistinguishable from your pre-op capacity to make color judgments. Assuming, of course, that her reasoning skill and her acquaintance with the color space are good enough.

Why do I say *nearly* indistinguishable? Because even though her post-op sight-based color judgments will be so good that in the normal course of things they will never lead anyone to suspect she underwent an ECI-op, they will not be *identical* to what they were pre-op. The reason is that in normal human color vision there are more just noticeable differences (JNDs) of shade between red and blue than there are between blue and green (Kay, 1999; Palmer, 1999). There will therefore be distinct shades $S1$ and $S2$ that pre-op Alice could just distinguish visually, but post-op she cannot; and there will be distinct shades $S3$ and $S4$ that pre-op she could not distinguish visually, but post-op she just can. This fact might make itself manifest if she is asked to make sight-based color similarity judgments about determinate shades that are very close to each other - the sort of judgment that one

typically does not make, but might be asked to make in a psychology lab.¹⁰

In the years after the ECI-op, Alice undertakes to make a few hundred sight-based color judgments a day. Over time, she finds that she can make subtler judgments faster, with less effort and sometimes without any conscious reasoning. By the time she is 50 (this is stage 3) this training regimen has long been abandoned. At stage 3, making sight-based color judgments is second nature to her. She makes them immediately, without any thinking or reasoning, and with effortless dexterity. In fact, the way she made sight-based color judgments before the ECI-op is identical, in all respects available to her consciousness, to the way she makes them now. Still, at stage 3 Alice remembers full well how things looked to her pre-op. She knows that fire trucks look to her the color fresh grass used to look to her. And as Taylor (1966) remarked, if she is asked to paint two pictures, one of how things looked to her pre-op and one of how they look to her post-op, she will do the following: In the post-op picture she will paint the grass with green paint. In the pre-op picture she will paint it with red paint.

At stage 4 Alice is 60 years old, and she has developed amnesia for the period up to age 50. In particular, she does not remember the following: (i) how things used to look to her before the ECI-op, (ii) that she so much as had the ECI-op, (iii) what the doctors taught her about how shades would look to her post-op, and (iv) that she ever arduously had to think before making color judgments. Alice also has no memories of episodes of remembering anything from the period up to age 50. Despite her amnesia, however, her color vision at stage 4 remains as it was at stage 2. There are still distinct shades $S1$ and $S2$ ($S3$ and $S4$) that pre-op she could (not) just distinguish, but post-op she cannot (can). Furthermore, sight-based color judgments are still second nature to her, and she makes them in the same way and with the same accuracy as in stage 3.

The intrapersonal spectrum inversion scenario can be used to argue against ER1, but only given two auxiliary definitions, and an auxiliary assumption. Definition 1: For any color or shade S , say e is an S -experience just in case, for some object x , e is veridical at w

¹⁰For a discussion of whether a behaviorally undetectable spectrum inversion scenario is empirically possible in humans, see Block (2007c); Broackes (2007); Hardin (1997); Harrison (1973); Hilbert and Kalderon (2000); Nida-Rümelin (1996); Palmer (1999).

iff x is S at w .¹¹ (For example, my experience of a certain red couch is a red-experience just in case it is veridical at w iff the couch is red at w .) Definition 2: e is a *color experience* just in case for some color or shade S , e is an S -experience. The auxiliary assumption is: *The intensions of color experiences determine the colors things look to the experiencer.*¹²

We can now argue as follows: Suppose Alice looks at a red couch in good light twice - once at stage 1, before her ECI-op, and again at stage 4, after she developed amnesia. Suppose that in both cases she directly comes to justifiably believe, of the couch, that it is red on the strength of a color experience.¹³ However, the couch looks opposite colors to Alice at stages 1 and 4. By the auxiliary assumption, it follows that at least one of the relevant color experiences is not a red-experience. So at some point Alice directly comes to justifiably believe something on the strength of an experience which the experience does not entail. So we have a counterexample to ER1.

Objection 1. The inverted spectrum scenario you describe is incoherent, inconceivable or otherwise fails to constitute an understandable scenario. The only way to make the scenario understandable is if we alter it so that Alice's color vision reinverts at either stage 3 or 4. At one of these stages it must become the case that for any shade S that an object can have, and for any lighting conditions C , things of shade S (in C) look to her the color they did pre-op (in C).

Reply. I take it to be obvious that stages 1 and 2 of the scenario are perfectly understandable, as we can certainly make sense of the idea that *we* will wake up tomorrow finding the world looking very strange indeed. But if the scenario is coherent at stage 2, why should anyone think that coherence can only be preserved if Alice's color vision reinverts at stage

¹¹This definition presupposes that the truth conditions of experiences are object-dependent. The presupposition is controversial. Those who reject it may instead say, for any color or shade S , that e is an S -experience just in case e is veridical at a centered world $\langle w, i \rangle$ iff for some object x appropriately related to i at w , x is S at w . Nothing in this chapter will hang on these matters.

¹²This assumption is part of the view of many phenomenal representationalists, including Alex Byrne and David Hilbert (1997), Fred Dretske (1995), Gilbert Harman (1990), Christopher Hill (2009), William Lycan (1996), John McDowell (1996; 1998) and Michael Tye (1995; 2000; 2009). The positions of David Chalmers (2010b; 2010a), Martin Davies (1992; 1997), Uriah Kriegel (2002), Adam Pautz (2006) and Sydney Shoemaker (1994) are more difficult to classify. Tyler Burge (2003b) and Ned Block (1990; 2007a; 2007b; 2007c) reject the assumption. (The term “phenomenal representationalists” was explained in footnote 1.)

¹³Some readers might hold that on the strength of experience one can only directly come to justifiably believe propositions of the form *SOMETHING-IS-F* and not propositions of the form *THAT-IS-F*. These readers may adjust the text accordingly. Nothing here hangs on this matter.

3?

According to Daniel Dennett (1988; 1991), this is because at stage 2 Alice would hesitate or stumble in making color judgments and this will serve as third person corroboration for her first person assertions that things look strangely colored to her. However, since at stage 3 her color judgments are perfectly normal, we have no way of determining whether her color vision is the same then as it was at stage 2 and she correctly remembers that it was different at stage 1, or whether her color vision reinverted and she is suffering from some memory malfunction. Our inability to decide between these alternatives shows that the scenario as described becomes incoherent at stage 3.

This objection merits a three pronged response. First, the objection presupposes that incoherence comes in the wake of an inability to decide between two alternative explanations of one phenomenon. This seems implausible. It is to confuse being asked a question to which we cannot find an answer with being asked no question at all. Second, it is simply not true that we have no way of determining whether the subject's color vision is the same at stages 2 and 3 or whether at stage 3 she is suffering from some memory malfunction. As Block (2007c) points out, there is no reason to doubt that we could decide between the two possibilities by empirical methods. E.g., perhaps exploration of the neural mechanisms involved in online visual experiences and iconic memory preservation might allow us to determine whether their functioning has changed in the period between stages 2 and 3. Third, given how the inverted spectrum scenario was described, we do not even have to resort to neural empirical methods. We *know* that at stage 3 Alice's memory was well functioning and that her color vision was as it was at stage 2. We know this because it is revealed by her capacities to make color discriminations. At stages 2 and 3 there are distinct shades $S1$ and $S2$ ($S3$ and $S4$) that she cannot (can) distinguish, but that she could (not) just distinguish at stage 1. What is the best explanation for this fact? Considering that normal human color vision is such that there are more just noticeable differences (JNDs) of shade between red and blue than there are between blue and green, the best explanation seems to be that her color vision was inverted at stage 2 and remained inverted at stage 3. So by reasoning to the best explanation we can come to know that at stage 3 Alice's color vision was as it was at stage 2.

What of the claim that the scenario becomes incoherent at stage 4? Once the coherence

of the scenario at stage 3 has been conceded, the objection to its coherence at stage 4 seems weak. For why should Alice's amnesia make the scenario incoherent? More specifically, why should Alice's amnesia incline us to doubt that her color vision remains as it was at stage 3? Furthermore, we *know* that at stage 4 Alice's color vision was as it was at stage 2. As we already noted, our grounds are that at stages 2, 3 and 4 there are distinct shades $S1$ and $S2$ ($S3$ and $S4$) that she cannot (can) distinguish, but that she could (not) just distinguish at stage 1.

Objection 2. Even though the inverted spectrum scenario you describe is coherent, it is nonetheless metaphysically impossible.

Reply. The first thing I think should be pointed out is that the objection is unmotivated. At stage 1 Alice is physically, functionally and behaviorally *different* from herself at stage 4. (Recall that the ECI-op involved physical changes to her nervous system, and that her spectrum inversion is behaviorally detectable.) The scenario is therefore simply not one of a physically, functionally or behaviorally undetectable inversion; and hence, it makes none of the assumptions that might have made its metaphysical possibility questionable. Of course, a supporter of ER1 that also accepts the auxiliary assumption driving the argument might be inclined to say that the scenario is impossible precisely because it can be used to argue against ER1. But such a reply is clearly question-begging in the current dialectical context.

The second and more important thing I think should be said in support of the possibility of the scenario is that we have some empirical evidence suggesting not only that red-green spectrum inversion cases are *empirically* possible for humans, but that they are *actual!* People with normal color vision (trichromats) have three kinds of pigments in their three cone types. People with a certain kind of red-green color blindness (protanopes) have a gene which causes their long-wavelength (L) cones to have the same pigment as their medium-wavelength (M) cones. People with another kind of red-green color blindness (deuteranopes) have a gene which causes their M-cones to have the same pigment as their L-cones. Both protanopes and deuteranopes cannot experience the red-green dimension of color space because the visual system codes the information mostly responsible for it as the difference between the outputs of the M- and L-cones. So how would things look colorwise to a person who had the genes responsible for *both* protanopia and deuteranopia? According to [Piantanida \(1974\)](#)

and [Boynton \(1979, p. 356-358\)](#), the person would have L-cones with the M-pigment and M-cones with the L-pigment, and so the difference between the outputs of the M- and L-cones would simply be systematically reversed in sign. The person would therefore not be color blind at all, but rather be a red-green inverted trichromat (or “pseudonormal”). It is estimated that 14 out of 10,000 males have pseudonormal vision. [This evidence was first brought to the philosophical community by [Martine Nida-Rümelin \(1996\)](#). For more on the debate surrounding this evidence, see [Block \(2007c\)](#); [Byrne and Hilbert \(2003\)](#); [Palmer \(1999\)](#).] Even if one is not persuaded that pseudonormals are actual, there is even stronger evidence that other kinds of intrapersonal color perception changes and interpersonal color perception differences are actual. In his [\(2007b\)](#), Block surveys three lines of empirical evidence that suggest that there are significant differences in the colors that objects look to normal perceivers, especially if they are of different sexes, races and ages. The lines of evidence are these: (i) There are significant differences between the visual systems of normals. In particular, the peak sensitivities of certain kinds of cones vary among normal subjects. (ii) There are significant differences in the way normals perform on color matching tests. In one kind of test, subjects are asked to make two halves of a screen match in color. One half of the screen is lit by a mixture of red and green light. The other half is lit by yellow or orange light. The subjects are asked to effect the match by controlling the intensities of the red and green lights. As it turns out, one normal perceiver will often typically effect a match using very different light intensity settings than another normal perceiver, especially if they are of different sexes. (iii) Normal subjects typically disagree greatly about which light wavelength shows up as unique green. This seems to be rather compelling evidence that the same thing (under the same lighting conditions) often *does* look different colors to different individuals (or to the same individual at different ages). If these kinds of physically and behaviorally detectable spectrum shifts are *actual*, it seems ludicrous to deny the metaphysical possibility of the described physically and behaviorally detectable spectrum inversion case.

Objection 3. A premise in your argument against ER1 is that both at stage 4 and at stage 1 Alice comes to believe, of the couch, that it is red. But at least at one of those stages the couch does not look red to her. So at least at one of those stages Alice does not come to believe, of the couch, that it is red.

Reply. This objection flies in the face of the profoundly historical and social nature of our beliefs, articulated in the seminal works of Tyler Burge (1979), Saul Kripke (1980) and Hilary Putnam (1975). One of the guiding insights of these works is that since there can be logical relations between my thoughts and your thoughts (e.g., they may contradict each other, in which case we would be in a disagreement), my thoughts and your thoughts had better be shareable. Now let us apply this insight to the case of color concepts. Suppose you meet with Alice twice - when she is looking at the red couch at stage 1 and when she is doing the same at stage 4. On both occasions Alice expresses what she comes to believe by asserting “That’s red”. On both occasions you, for some reason, believe differently. You express your belief by asserting “No, that isn’t red”. So on both occasions you and Alice would be in a genuine disagreement, since on both occasions you would have contradictory beliefs. But since nothing strange is going on with *your* visual system, there is no question about what *you* believe on either of the two occasions: You believe, of the couch, that it is *not* red. So it must be the case that Alice (both at stage 1 and at stage 4) comes to believe of the couch, that it *is* red.

An objection will now be made: On one of these occasions you and Alice do not really disagree. Perhaps you only *seem* to disagree about the color of the couch because the sentences you utter sound incompatible. Couldn’t it be that you are just using the English word “red” differently, to refer to different things? No. What shows that the disagreements are both genuine is the fact that on both occasions you two are rationally responsive to the same considerations about the color of the couch. For example, you could persuade one another by calling on further eye witnesses, by presenting each other with results on color matching studies done on the couch, by showing each other scientific figures regarding the couch’s spectral surface reflectance function, etc. Another fact that shows the disagreements to be genuine is this: If one side persuaded another that the couch is (not) red, then the persuaded side would *not* say, “my views about the couch were *true* before, only I should have expressed them by applying a different color term to the couch”. Instead, the persuaded party would say “my view about the couch was *false*”.

Objection 4. A premise in your argument against ER1 is that both at stage 4 and at stage 1 Alice comes to justifiably believe, of the couch, that it is red. But at least at one of

those stages the couch does not look red to her. So at least at one of those stages Alice does not come to justifiably believe, of the couch, that it is red.

Reply. I will give three replies to this objection:

(i) I take it to be deeply intuitive that Alice's pre-op experience-based color judgments, and her pre-op judgments about the colors things look her, are justifiably made. Why? Well, largely for selfish reasons: I take it to be obvious that my own experience-based color judgments, as well as my judgments about the colors things look to me, are typically justifiably made. But there is no relevant difference between pre-op Alice and myself. In fact, the entire thought experiment could be run with me as its hero, instead of Alice. So it seems very hard for me to deny Alice the same status. And since the reader can go through the same selfish line of reasoning too, the reader should agree. Now: If Alice's judgments about the colors things look her are justifiably made at stage 1, how could her experience-based color judgments fail to be justifiably made at stage 2? After all, at stage 2 her experience-based color judgments are *logically inferred* from (a) her judgments about the colors things look to her - which are justifiably made because she makes them as she made them at stage 1, and (b) her *knowledge* (which she acquired through impeccable training), for any shade S and lighting conditions C , what shade things of shade S (in C) will look to her post-op. So at stage 2, Alice's experience-based color judgments are justifiably made as well. But, if these judgments are justifiably made at stage 2, how can they fail to be similarly justifiably made at stages 3 and 4? What sets stage 3 apart from stage 2 is that by stage 3 Alice has become an expert judge of color. The acquisition of expertise, however, cannot make the judgments that issue from it *worse* than those that issued from a novice. Similarly, what sets stage 4 apart from stage 3 is that Alice suffers from amnesia at stage 4. But the amnesia has no impact on Alice's expertise at making color judgments, just as it has no impact on any of Alice's other areas of expertise - making experience-based judgments about the locations of things, or about the identities of familiar people, for example. Since it has no impact on the expertise, it has no impact on the justifiability of making the judgments that issue from it.

(ii) There are a number of theoretical considerations supporting the view that Alice's experience-based color judgments (both at stage 1 and at stage 4) are justifiably made: (a) At stage 1, we may assume, Alice's color judgments are highly reliable. And it is part of

the description of the case that at stages 2-4 Alice's capacity to make color judgments is nearly indistinguishable from her stage 1 capacity to make them. So at all stages, Alice's color judgments are highly reliable. A reliabilist will therefore be forced to admit that at all stages, Alice's color judgments are justifiably made. (b) It follows from the description of the case that at stage 4 the way Alice makes sight-based color judgments is identical, in all respects available to her consciousness, to the way she made them at stage 1. So at both stages Alice makes color judgments that are experience-based, free of any inclinations to believe otherwise, free from contradictory beliefs, free from misgivings, and free of any other defeaters. So the internalist epistemologist should have no reason to object to the justifiability of forming these judgments. (c) Both at stages 1 and 4 Alice could offer the standard sorts of justifications we all give for our experience-based color judgments. In particular, if asked why she believes, of the couch, that it is red, Alice could honestly and truthfully reply by uttering (on either occasion) "because I can tell a red thing when I see it", or "because I see that it is red". This is a further reason for the internalist epistemologist to accept that Alice's experience-based color judgments are typically justifiably made.

(iii) In response to a previous objection we reviewed empirical evidence that certain kinds of spectrum shifts and inversions are real, and even common. But if these cases are real, then we are in the following situation: Even though actual people often agree in their experience-based color judgments, that same couch sometimes looks very different shades to them. But even when we know this fact, we are not moved in the slightest to doubt the justifiability of each other's color judgments. So, in fact, we do not take the color that something looks to a person to be a deciding factor in estimating the justifiability of that person's experience-based color judgments. And it seems highly implausible and deeply unmotivated to suppose, despite our actual inclinations, that the color that something looks to a person *is* a deciding factor in estimating the justifiability of that person's experience-based color judgments.

Objection 5. A premise in your argument against ER1 is that both at stage 4 and at stage 1 Alice directly comes to believe, of the couch, that it is red. But surely at stage 4 Alice's experience-based color judgments are only indirectly made.

Reply. Before turning to the objection proper, let us ask ourselves if a supporter of ER can afford to claim that Alice's experience-based color judgments were *indirectly* made at

stage 1. The answer must be “no”. The normal person’s experience-based judgments about the distribution of sensible properties (color, shape, location, motion, texture, etc.) are the ER-supporter’s paradigm examples of directly formed beliefs.¹⁴ If these judgments are not directly made, then there is no plausibility in the idea that *any* judgments are directly made. So if the notion of a directly made experience-based judgment is not to become an empty one, it had better be the case that the normal person’s experience-based color judgments are directly made. Since at stage 1, before her ECI-op, Alice just *is* a normal person, her experience-based color judgments are directly made as well.

We may now turn to the objection proper. Our strategy will be to argue that there is no relevant difference between the way Alice forms experience-based color beliefs at stage 1 and the way she forms them at stage 4. Since at stage 1 Alice’s way of forming these judgments is direct, it must also be direct at stage 4.

In a way, it is part of the description of the scenario that there is no relevant difference between the way Alice forms experience-based color beliefs at stage 1 and the way she forms them at stage 4. After all, it is part of the scenario that at stage 4 the way Alice makes sight-based color judgments is identical, in all respects available to her consciousness, to the way she made them at stage 1. Still, it will be objected that there are *unconscious* differences between the stage-1- and the stage-4-way of forming these judgments, which are sufficient to render the stage 4 judgments indirectly made. What would these unconscious differences be? (a) Perhaps it will be claimed that at stage 1 Alice’s judgments are directly made because they are not inferentially made, whereas at stage 4 they are indirectly made because they are *unconsciously* inferentially made. This, however, cannot be right. Alice suffers amnesia at stage 4, and so she simply loses the beliefs and the other extra-experiential resources which would be necessary for her if she indeed came to her color judgments inferentially. Hence, Alice’s stage 4 color judgments are neither consciously inferentially made nor unconsciously inferentially made. (b) Perhaps it will be claimed that at stage 1 Alice’s judgments are directly made because they are of innate origins, whereas at stage 4 they are indirectly made because they are more profoundly the result of special training. This is empirically false. Alice’s stage 4 expertise in making experience-based color judgments is the result

¹⁴See, e.g., Burge (2003a); McDowell (2009); Peacocke (1992, 2004)

not of what the doctors taught her at stage 1 (which she forgot), but of the training she undertook between stages 2 and 3. Through this training she acquired the capacity to respond to experiences in which something looked a certain color to her by coming to justifiably believe, of the thing, that it is a certain color. But precisely the same is true of Alice at stage 1, or for that matter, of you and me. As every parent knows, it takes significant training before a human child can respond to experiences in which something looks a certain color to her by coming to justifiably believe, of the thing, that it is a certain color. So at stage 1 Alice's color judgments are no more innate than they are at stage 4.

(c) Perhaps it will be claimed that at stages 1 and 4 Alice's judgments are directly made because they are facilitated through the operation of certain computational modules or neural mechanisms, whereas at stage 4 they are indirect because they are facilitated by different modules or mechanisms. This claim imposes a philosophical preconception on an open empirical question. Which modules or mechanisms would facilitate Alice's judgments at the various stages is something that should be decided by empirical inquiry, and not from the armchair. Furthermore, since the scenario only needs to be *metaphysically* possible, it seems that it is possible to nip the present objection in the bud by simply *stipulating* that the same modules or mechanisms facilitate Alice's color judgments in stages 1, 3 and 4.

(d) Finally, perhaps it will be claimed that at stage 1 Alice's judgments are directly made because the color propositions Alice comes to believe by making them are entailed by Alice's relevant color experiences, whereas at stage 4 Alice's judgments are indirectly made because the color propositions Alice comes to believe by making them are not entailed by Alice's relevant color experiences. My response is that I am willing to accept that at stage 1 Alice comes to believe (on the strength of experience) color propositions that are entailed by her experiences. I am also willing to accept that at stage 4 Alice comes to believe (on the strength of experience) color propositions that are *not* entailed by her experiences. I am willing to accept both claims since it is part of my own argument *against* ER1 that either at stage 1 or at stage 4 Alice comes to believe a color proposition that is not entailed by the relevant experience. What I object to is the exploitation of the idea that Alice's stage 4 color beliefs are not entailed by the relevant experiences to argue that Alice comes to form these beliefs indirectly. Making this move requires the following claim to be a *stipulation*: One can directly form the belief

that p on the strength of an experience e only if e entails that p . Since this stipulation *entails* ER1, making it will make ER1 true by stipulation. But ER1 was meant to be an *informative* thesis. ER1 was supposed to illuminate a familiar phenomenon - a person's having the ability to directly come to justifiably believe something on the strength of an experience - by placing an informative constraint on it, namely, the obtaining of certain semantic relations between the experience and the belief. Such illumination will not be forthcoming if ER1 is to be true by stipulation, since one cannot illuminate a phenomenon by saying that it meets some constraint by stipulation. So if ER1 is to be true by stipulation, it will fail to be illuminating.

Having answered these five objections, let us conclude that ER1 (if potentially illuminating) is false, at least given the auxiliary assumption that the intensions of color experiences determine the colors things look to the experiencer. In the rest of the chapter I will attempt to expand on this argument and show that ER is false even without the auxiliary assumption.

2.3 EXPERIENCE-FIRST ACCOUNTS

Some supporters of ER—e.g., McDowell (1996; 1998)—accept the auxiliary assumption that the intensions of color experiences determine the colors things look to the experiencer. Other supporters however—e.g., Burge (2003b)—reject it. It is therefore open for them to claim that at stage 1 and 4 Alice underwent a red-experience (as defined above) when she looked at the red couch. And if Alice underwent a red-experience on both occasions, the previous section's argument against ER1 breaks down. In this section I will begin to remedy this breakdown.

Theories of content, for all their diversity, have a standard structural feature - they assign contents to *token* mental states by first assigning it to mental state *types*.¹⁵ In other words, these theories first put forward a condition a mental state *type* has to satisfy to have a certain content (e.g., that tokens of that type be normally, or regularly, caused in a certain way, or have a certain kind of function). Then they add that *token* mental states have the content of

¹⁵To be more accurate, only theories of non-indexical content have the described structural feature. Theories of content for indexical expressions, such as “I”, “here” or “now”, function differently. This complication will not be relevant here.

the type they belongs to. This structural feature has many advantages, but it presupposes that some non-semantic criterion of individuation induces a classification of mental state tokens into types. Importantly, this criterion of individuation can be changed independently of the rest of the theory of content, with the effect that the theory assigns different contents to different token mental states.

Visual experiences of color can be non-semantically classified into types in broadly three ways: (i) by the colors things look to the experiencer who has them, (ii) by the neural states that realize (or proximally cause) them, or (iii) by the *narrow* functional roles they play.¹⁶ Let us call “experience-first accounts” those accounts that (at least for the purposes of theories of content) classify visual experiences of color into types in one of the first two ways. Let us call “function-first accounts” those accounts that classify visual experiences of color into types in the third way. On experience-first accounts, Alice’s experiences of the red couch at stages 1 and 4 are not of the same type.¹⁷ On function-first accounts, they are.¹⁸

This section will be devoted to showing that ER1 must be rejected by anyone who accepts an experience-first account. The next section will complete the argument against ER by showing that ER2 must be rejected by anyone who accepts a function-first account.

Consider the scenario *Bob*: At stage 1, Bob is 18 and has normal color vision. After the sun sets and Bob falls asleep, Bob is unknowingly taken to a hospital, where his kidnappers perform a red-green ECI-op on him. The chip placed in his optic nerve operates as it did in the *Alice* scenario, and the operation has the usual effect. But this is not the only operation

¹⁶Can we not also classify visual experiences into types by their *wide* functional roles? On pain of having no solution to the disjunction problem, we cannot. In the case of experiences, the disjunction problem is that if experiences are to ever count as hallucinations or illusions, it must be possible for them to have the same content as their non-hallucinatory and non-illusory counterparts. But illusions and hallucinations often *fail* to be caused by the appropriate causes or in the usual ways. They also often fail to cause the appropriate behaviors or to do so in the usual ways. Therefore, if we classify visual experiences into types by their *wide* functional roles, this will have the result that illusions and hallucinations on the one hand and their non-hallucinatory and non-illusory counterparts on the other will fall under distinct functional types. If they fall under distinct functional types, there is every reason to expect that they will have *distinct* contents, and worse, that illusions and hallucinations will be systematically *veridical* when assigned these distinct contents. The disjunction problem will then be unsolvable.

¹⁷Experience-first accounts have this result for two reasons. First, in having the two experiences, the couch looks distinct colors to Alice. Second, the chip in Alice’s brain ensures that the two experiences involve different kinds of neural states.

¹⁸Function-first accounts have this result because the two experiences play the same narrow functional role in Alice - they are caused by the same pattern of retinal input, they involve (nearly) the same conceptual classifications, and they cause (nearly) the same behavioral output.

the kidnappers perform on Bob. They also perform a red-green “Conceptual Color spectrum Inversion” operation (a “CCI-op”, for short) on him. The operation involves placing a chip in the area of Bob’s brain that is responsible for the conceptual color classifications Bob makes *directly* on the strength of his color experiences. The chip performs the following function: When it is on, Bob classifies (directly and on the strength of experience) things that look to him the colors that green and greenish things (in C) looked to him pre-op, under the concepts that pre-op he used to apply to red and reddish things (in C). Similarly, Bob classifies (directly and on the strength of experience) things that look to him the colors that red and reddish things (in C) looked to him pre-op, under the concepts that pre-op he used to apply to green and greenish things (in C). Finally, suppose that both chips are usually off, and that when they are off they have no impact on Bob whatsoever, so that things with Bob are just as they would be if he had undergone neither operation. In fact, suppose the chips are designed to turn on together randomly, for one minute only, and then to turn back off forever. Let stage 2 be that one minute period at which they are on.

What happens at stage 2? The joint effect of the operation of both chips is that during stage 2 Bob makes the same conceptual color classifications that he would have made if he had undergone neither operation. To see this, consider an example: Suppose that at stage 2 Bob looks at a red object in good light. The ECI-chip makes that object look to Bob the color that green things in good light looked to him pre-op. But the CCI-chip then makes Bob classify things that look that color to him under the concepts that pre-op he used to apply to *red* things in good light. So Bob (directly and on the strength of experience) comes to believe of the object that it is red.¹⁹

The *Bob* scenario is nothing more than a time-compressed version of the *Alice* scenario. At stage 2 Bob is in the same situation as Alice is in at stage 4. The only differences are these: (i) Where Alice had to laboriously and intently acquire her stage 4 expertise at making experience-based color judgments, Bob is given his stage 2 expertise at making experience-based color judgments effortlessly, through the workings of a microchip. (ii) Where Alice

¹⁹*Caveat*: What was just claimed is not perfectly accurate. If at stage 2 Bob is asked to make sight-based color similarity judgments about determinate shades that are very close to each other, his judgments may not be identical to those he would have made if he if he had undergone neither operation. As in *Alice*, the reason is that asymmetries of human color perception mean that there are distinct shades $S1$ and $S2$ ($S3$ and $S4$) such that pre-op on Bob could (not) just distinguish, but when the ECI-chip is on he cannot (can).

had to suffer amnesia to rid herself of the capacity to make color judgments by inferential means, Bob never comes to have this inferential capacity in the first place.

Since at stage 2 Bob is in the same situation as Alice in stage 4, it cannot be objected that Bob's situation at stage 2 is incoherent or metaphysically impossible. Neither can it be claimed that at stage 2 Bob cannot come to have the same color beliefs that he would have come to have had he undergone neither operation. The reasons for this are the same as those offered in response to objection 3 of the previous section. Can it be claimed that at stage 2 Bob's experience-based color judgments are not justifiably made? I see no hope for this. At stage 2 Bob is in the same situation as Alice in stage 4. Since at stage 4 Alice's experience-based color judgments are justifiably made, so are Bob's judgments at stage 2. The justified formation of Bob's stage 2 beliefs is furthermore supported by the considerations rehearsed in parts (ii) and (iii) of the response to objection 4 of the previous section.²⁰ Finally, it is part of the description of the case that Bob's experience-based color judgments are made directly at stage 2. The legitimacy of building this feature into the description of the case should be obvious given the considerations rehearsed in response to objection 5 of the previous section.

²⁰ *Objection.* But at the instant that stage 2 begins, the colors things look to Bob will change radically. Surely Bob will then believe that the colors things look to him have changed. And having come to believe this, he will have a defeating consideration against making any experience-based color judgments. If he proceeds to make such judgments despite this defeater, these judgments would consequently not be justifiably made. Similarly, suppose at stage 2 Bob accesses a pictorial memory of a familiar object (a fire truck, say). Comparing the color fire trucks look to him at stage 2 with the color they visually appear in the pictorial memory, Bob will again come to believe that the color things look to him has changed. So he will again have a defeater against making any experience-based color judgments.

Reply. The first part of the objection relies on the idea that Bob will notice the change in the colors that things look to him. But Bob will precisely *not* notice this change. To see why, suppose Bob is looking at a fire truck just before stage 2 begins. On the strength of his color experience, Bob directly comes to classify the truck under various color concepts (e.g., the concepts "red" and "looks red"). Now stage 2 begins, and both the ECI- and the CCI-chips turn on. Given the way they work, it follows that Bob will *continue* to directly classify the truck under the very same color concepts. Since Bob's classifications remain the same once stage 2 begins, Bob will *not* come to believe that the colors things look to him have changed.

The second part of the objection relies on the idea that once stage 2 has begun, Bob will notice the difference between the colors that things look to him and the colors they visually appear when Bob accesses pre-stage-2 pictorial memories of them. To circumvent this objection we need only suppose that the ECI-chip (once on) impacts not only the colors things look to Bob, but also the colors they visually appear when Bob accesses pre-stage-2 pictorial memories of them. To see the force of this supposition, suppose Bob accesses (at stage 2) a pictorial memory of a fire truck. The supposition means that it will visually appear to him (in the memory) the color fresh grass used to look to him pre-op. Consequently, it will visually appear to him (in the memory) the color fire trucks look to him at stage 2. A result of this is that Bob will classificatorily respond to the memory of the fire truck in the same way that he would classificatorily respond to an on-line experience of the fire truck. So once again Bob will *not* come to believe that the colors things look to him have changed.

Here then is the argument showing that ER1 must be rejected by anyone who accepts an experience-first account: Suppose that at stage 2 Bob looks at a red couch in good light. Suppose that on the strength of his experience he directly comes to justifiably believe, of the couch, that it is red. We will soon see that on any plausible theory of content complemented by an experience-first account, Bob's stage 2 experience is a green-experience (in our technical sense). It follows that at stage 2, Bob directly comes to justifiably believe, of the couch, that it is red on the strength of a green-experience. So, at stage 2, he directly comes to justifiably believe something on the strength of an experience which the experience does not entail. So we have a counterexample to ER1.

To complete the argument it remains to show that on any plausible theory of content complemented by an experience-first account, Bob's stage-2 experience is not a red-experience. This will be accomplished presently by reviewing the consequences of five leading theories of content complemented by an experience-first account. To begin, note that Bob's stage 2 experience is one in which something looks to him the color green things used to look to him in good light pre-op, and that it is realized (proximally caused) by the same kind of neural state as those pre-op experiences. Therefore, on experience-first accounts Bob's stage 2 experience is of the same type as his stage 1 experiences of green things (seen in good light). Call experiences of this type "GEs". Now consider the five theories of content:

1. *Causal Theories.* According to these theories, the intension of an experience e is determined by the states-of-affairs that normally, or regularly, cause experiences of e 's type. Bob's stage 2 experience is a GE. Since stage 2 lasts only a minute, the normal, or regular, causes of GEs do not change. These causes are green things. So Bob's stage 2 experience is not a red-experience, but a green-experience.

2. *Teleosemantic Theories.* According to these theories [Burge (1986; 2003a; 2010), Dretske (1995), Millikan (2004), Peacocke (2004)], the intension of an experience e is determined by the states of affairs that experiences of e 's type have the evolutionarily selected function of reliably indicating. Bob's stage 2 experience is a GE. Since Bob's operations have no impact on his evolutionary history, the evolutionarily selected function of GEs does not change. This function is to reliably indicate green things. Therefore, Bob's stage 2 experience is not a red-experience, but a green-experience.

3. *Asymmetric Dependence Theories.* According to these theories [Fodor (1992)], the intension of an experience e is determined (roughly) by the states-of-affairs-type X such that (i) X tokens lawfully cause experiences of e 's type, and (ii) for any states-of-affairs-type Y distinct from X such that Y tokens lawfully cause experiences of e 's type: (a) if Y tokens did not lawfully cause experiences of e 's type, X tokens still would, and (b) if X tokens did not lawfully cause experiences of e 's, Y tokens would not either. Bob's stage 2 experience is a GE. Since stage 2 lasts only a minute, there is no change to what states-of-affairs *lawfully* cause GEs. In fact, GEs asymmetrically depend on things being green. Therefore, Bob's stage 2 experience is not a red-experience, but a green-experience.

4. *Optimal Conditions Theories.* According to these theories [Tye (1995; 2000)], the intension of an experience e is determined (roughly) by the states-of-affairs-type X such that if optimal conditions were to obtain, the experiencer would have an experience of e 's type iff, and because, it is the case that an X -type state of affairs obtains. Bob's stage 2 experience is a GE. Under optimal conditions, Bob has GEs iff, and because, something is green. So his stage 2 experience is a green-experience.

5. *Causal, Conceptual or Inferential Role Theories.* This family of theories [Block (1986; 1987), Harman (1982; 1987)], which grows out of the functionalist tradition, is always complemented by function-first accounts. Still, as idiosyncratic as it is, it is possible to describe a hypothetical causal, conceptual or inferential role theory that is complemented by an experience-first account. On this mongrel theory, the intension of an experience e is determined by the *wide* causal, conceptual or inferential role played by experiences of e 's type—individuated experience-firstly. Bob's stage 2 experience is a GE. Since stage 2 lasts only a minute, the *wide* causal, conceptual or inferential role GEs play does not change. This wide role involves being caused by green things, causing beliefs to the effect that something is green, and leading to appropriate behaviors. Therefore, Bob's stage 2 experience is not a red-experience, but a green-experience.

Conclusion: When experiences are typed in experience-first ways, plausible theories assign them contents independently of the content of the beliefs these experiences may give rise to. The *Bob* scenario exploits this independence. Examined through experience-first eyes, the scenario is a case where an experience is assigned content which does not entail

the content of the belief the experience gives rise to. Since the belief is also had directly and justifiably, the scenario is a counterexample to ER1.

2.4 FUNCTION-FIRST ACCOUNTS

The reader might well have noticed that the argument of the previous section falls flat if the various theories of content are complemented not with an experience-first account but with a function-first account. To take the most natural example of this, suppose we complement an inferential role theory with a function-first account. This marriage leads to the following simple theory: The intension of an experience e is determined by the *wide* causal, conceptual or inferential role played by experiences of e 's *narrow* functional type. Applying this theory to the *Bob* scenario has the following result: Since Bob's stage 2 experience of the red couch is caused by a red object, and leads to a judgment, of the couch, that it is red (which, in turn, leads to appropriate behaviors), it plays (almost exactly) both the wide and the narrow causal, conceptual or inferential role that was played by Bob's stage 1 experience of the red couch. Hence, both these experiences are red-experiences. So the argument against ER1 fails.

Turning to function-first accounts will not, however, save ER. The turn will only make ER2 vulnerable instead. [Reminder: ER2 says that if (i) e entails that p , and (ii) one directly comes to believe that p on e 's strength (where one is rational and has no defeaters), *then* one justifiably believes that p .] A counterexample to ER2 is given by a third thought experiment - the *Caitlin* scenario:

At stage 1, Caitlin is 18 and has normal color vision. After the sun sets she is kidnapped and taken to a hospital, where her kidnappers perform a red-green "Functional Color spectrum Inversion" operation (an "FCI-op", for short) on her. The operation consists in the installation of red-green color inverting lenses on Caitlin's eyes. The lenses can be turned on and off by turning a lens-controlling FCI-device on and off. When the lenses are off, they have no effects on Caitlin's vision. The FCI-device—and the controlled lenses as well—are usually off. They are all designed to turn on for one minute only, and then to turn back

off forever. Let stage 2 be that one minute period at which the device and lenses are on. Suppose that stage 2 begins when Caitlin first opens her eyes after a long sleep, that during stage 2 Caitlin never looks at anything whose color is familiar to her, and that immediately after stage 2 she falls asleep again.

What happens to Caitlin during stage 2? Suppose she looks at a red object in good light. The inverting lenses then proximally stimulate her eyes with green illumination. Caitlin therefore directly judges the experienced red object to be green, and is disposed to behave accordingly. Therefore, when Caitlin looks at a red object in good light at stage 2, she has an experience that plays the *narrow* functional role of her stage 1 experiences of green things in good light. At the same time, since stage 2 begins when Caitlin first opens her eyes, and since she never looks at anything whose color is familiar to her at stage 2, Caitlin has no reason whatsoever to be cautious about her stage 2 experience-based color judgments. So the effect of the FCI-op is that at stage 2 Caitlin completely loses her ability to tell the color of red and green things by sight.

It cannot be objected that Caitlin's situation at stage 2 is incoherent or metaphysically impossible. After all, we know full well that *we* would be in Caitlin's exact position if we were to undergo an FCI-op. Neither can it be objected that at stage 2 Caitlin either loses command of her *RED* and *GREEN* concepts, or that the semantics of these concepts are somehow changed. The reasons are the same as those offered in response to objection 3 of section 2.2. Furthermore, the legitimacy of supposing that Caitlin's experience-based color judgments are made directly is defended in the considerations rehearsed in response to objection 5 of section 2.2. Finally, since Caitlin has no reason whatsoever to be cautious about her experience-based color judgments, she has no available defeaters for any of the experience-based color beliefs she forms during stage 2.

We now have in place most of the ingredients we need to establish a counterexample to ER2. Only two pieces are missing from the puzzle. First, we need to see that on function-first accounts, when Caitlin looks at a red object in good light at stage 2, she has a *green*-experience (in our technical sense). The second thing we need to see is that, at stage 2, when Caitlin directly comes to believe, of a red object, that it is green (on the strength of her experiences), she does not justifiably so believe. Let us take these tasks in turn.

At stage 2, when Caitlin looks at a *red* object in good light, her color inverting lenses proximally stimulate her eyes with green illumination, on the strength of which she judges the experienced object to be green and is disposed to behave accordingly. Therefore, on function-first accounts Caitlin's experience is of the same *narrow* functional type as her stage 1 experiences of green things in good light. Call experiences of this type "GFs". Now consider the five theories of content:

1. *Causal Theories*. At stage 2, when Caitlin looks at a red object in good light she has a GF experience. Since stage 2 lasts only a minute, the normal, or regular, causes of GFs do not change. These causes are green things. So Caitlin's experience is a green-experience.

2. *Teleosemantic Theories*. At stage 2, when Caitlin looks at a red object in good light she has a GF experience. Since Caitlin's operation has no impact on her evolutionary history, the indication function of GFs does not change. This function is to indicate green things. So Caitlin's experience is a green-experience.

3. *Asymmetric Dependence Theories*. At stage 2, when Caitlin looks at a red object in good light she has a GF experience. Since stage 2 lasts only a minute, there is no change to what GFs asymmetrically depend on. They asymmetrically depend on things being green. So Caitlin's experience is a green-experience.

4. *Optimal Conditions Theories*. At stage 2, when Caitlin looks at a red object in good light she has a GF experience. Since Caitlin's operation has no impact on what counts as "optimal conditions" for her, the optimal causes of GFs remain clear: Under optimal conditions, Caitlin has GF experiences iff, and because, something is green. So Caitlin's experience is a green-experience.

5. *Causal, Conceptual or Inferential Role Theories*. At stage 2, when Caitlin looks at a red object in good light she has a GF experience. Since stage 2 lasts only a minute, the *wide* causal, conceptual or inferential role GFs play does not change. This wide role involves being caused by green things, causing beliefs to the effect that something is green, and leading to appropriate behaviors. So Caitlin's experience is a green-experience.

It now only remains to show that, at stage 2, when Caitlin directly comes to believe, of a red object, that it is green (on the strength of her experiences), she does not justifiably so believe. The reason for this is straightforward: The activation of the FCI-device robs

Caitlin of her expertise, or capability, to tell the colors of red (and green) things by sight. Although she is confronted with perfectly ordinary colors, her experience-based directly made color judgments go terribly astray. So at stage 2 Caitlin is simply out of touch with color reality. The relevant judgments she makes are radically unreliable and are formed in radically unreliable ways. Since Caitlin's directly made experience-based color judgments are contra-posed to the real distribution of redness and greenness, they cannot be justifiably made.

If you are unsure of this, consider the following: At stage 1 Caitlin is perfectly normal, and directly makes experience-based color judgments perfectly justifiably. If at stage 2 she also directly makes experience-based color judgments perfectly justifiably, then it follows that Caitlin can directly make inconsistent experience-based judgments about a single color reality equally justifiably. And this is surely absurd.

We are now ready to see why function-first accounts make ER2 false: Suppose Caitlin looks at a red couch in good light at stage 2 and directly forms, on the strength of her experience, the belief, of the couch, that it is green. We just saw that Caitlin's experience is a green-experience. We also noted that Caitlin is generally rational and equipped with the relevant color concepts. Finally, Caitlin has no defeaters for believing, of the couch, that it is green. From all that we said it would follow, if ER2 were true, that Caitlin *justifiably* believes what she does. But we also argued that this is something that Caitlin precisely does not do. So ER2 is false.

Objection. You are insisting that Caitlin does not justifiably believe, of the couch, that it is green. But this insistence is out of place. Caitlin believes what she does perfectly justifiably, since (i) she believes it in the absence of defeaters, and (ii) her belief is entailed by her grounding experience. This is all there is to it.

Reply. I should begin by making clear that I agree that Caitlin's belief is not entirely without virtue. The epistemic virtue of Caitlin's belief is that the belief is *rationally* had. To say that the belief is had rationally is to say that Caitlin forms it non-recklessly, i.e., while demonstrating sufficient care for the belief's veracity. This is allowed for by the fact that Caitlin is misled on the strength of an experience that is *not* appreciably misleading to her, and while having no defeaters for her belief.

Having granted that Caitlin's belief is rationally had, I need not resist the thought that the belief is so had because it is had in the absence of defeaters and on the strength of an entailing experience.²¹ Still, I do mean to resist the thought that the same considerations show that the belief is justifiably had.

If the notion of justification is to do serious epistemic work, it must be understood against the background of the notion of knowledge. At a minimum, this suggests the following: A belief cannot be had justifiably if it is formed in a way that *guarantees* that it and similarly formed beliefs would fall short of knowledge. This is precisely the problem with Caitlin's belief. The operation of the color-inverting lenses at stage 2 has the effect that Caitlin's directly formed experience-based belief, along with similar stage 2 beliefs about the redness and greenness of things, are guaranteed to fall short of knowledge. Therefore, there is no hope for the view that in addition to being had rationally, Caitlin's belief is also had justifiably.

If the reader still feels gripped by the position expressed in the objection, she might consider this: On the present conception of contents, Caitlin's experience indeed entails the belief, of the couch, that it is green. But the fact that the belief is entailed by the grounding experience contributes nothing to the belief's justification. It contributes nothing because the intension of the grounding experience itself contributes nothing to Caitlin's chances of knowing the color of the couch. In fact, since the grounding experience is a green-experience, its intension is nothing if not misleading. We should therefore say, on the one hand, that Caitlin's belief is had rationally (potentially because it is entailed by its grounding experience), and on the other hand, that Caitlin's belief is not had justifiably (because the grounding experience profoundly undermines Caitlin's chances of knowing the color of the couch).

Objection. I still think Caitlin's stage 2 judgment about the couch was justifiably made! The color inverting lenses made the couch look to Caitlin the color green things in good light looked to her at stage 1. Caitlin's belief-forming process is to respond to experiences in which something looks that color to her by directly coming to believe, of the thing, that it is green. And since this belief forming process is perfectly reliable in normal settings (e.g., at stage 1), all beliefs formed by it are justifiably had. This includes Caitlin's stage 2 judgment about the

²¹I do have reservations about this, but I will not raise them here.

couch.

Reply. I think that we must reject the idea that what decides if a belief is justifiably had is whether it is formed by a process that is reliable in one's normal environment. To see why, think back to the *Bob* scenario. At stage 2 of that scenario we saw that Bob *justifiably* judged the red couch to be *red* (on the strength of experience). He did so despite the fact that the inputs to his belief forming processes would normally (e.g., at stage 1) be caused by *green* things. So at stage 2 Bob justifiably had a belief, despite the fact that at play were belief forming processes that would have been terribly *unreliable* in normal settings. So the reliability of a belief forming process in normal settings is not a *necessary* condition for the beliefs that issue from it to be had justifiably.

Still, someone might think that the reliability of a belief forming process in normal settings is a *sufficient* condition for the beliefs that issue from it to be had justifiably. And someone might further think that it is this fact that makes Caitlin's stage 2 judgment about the couch a justifiably made one.

The last position, however, has very little to recommend itself. First, it is unclear why the reliability of a belief-forming process should be sufficient but not necessary for justification. Second, any intuitive inclination one might feel towards the position can be explained away as the intuition that Caitlin's belief is rationally had. Nothing about justifiability follows from this. Third, the position violates the intuitive constraint on justification, according to which a belief cannot be had justifiably if it is formed in a way that guarantees that it and similarly formed beliefs would fall short of knowledge. The position violates this constraint because a belief formed by a process that is reliable in one's normal environment can also be guaranteed not to amount to knowledge when the belief-forming process is put to work in epistemically unfavorable circumstances, which render it unreliable. Beliefs formed by such an unreliable process cannot amount to knowledge and are therefore not justifiably had.

2.5 A LAST LINE OF DEFENSE

The arguments of the last three sections show that ER must be rejected. Together the arguments show that on any plausible way to assign contents to experiences, either ER1 or ER2 turns out to be false. So it seems that the view cannot be upheld. Still, before we wash our hands of ER, we should consider one way in which it might be defended.

The defender I am imagining is someone who, like Pryor (2000), takes ER to be true on intuitive grounds alone. Such a defender might put his view thus: “You are only managing to find problems with ER because we don’t have a good theory for assigning contents to experiences. But we can reasonably accept ER even if we don’t have a good theory of content. What makes this reasonable is that we have a pre-theoretical and intuitive conception of the contents of experiences which vindicates ER. On this conception, the content of a visual experience is just what, in having the experience, seems to the experiencer to be the case. For example, when you see a red couch, it seems to you that it is red. (Equivalently: When you see a red couch, the couch seems to you to be red.) Since the idea of something’s seeming to one to be the case is perfectly intuitive, I don’t need to say anything more about how to assign contents to visual experiences. I can just rest assured that both ER1 and ER2 are true on this intuitive conception of contents.”

Heroic as this defense is, it has no chance of succeeding. It is defeated by the *Caitlin* scenario. Consider: When Caitlin looks at the red couch at stage 2, she has an experience that plays the same narrow functional role as her stage 1 experiences of green things (seen in good light). Experiences with that functional role cause things to seem to be green to Caitlin. So, in having the stage 2 experience of the couch, the couch seems to Caitlin to be green. So, on the heroic defense’s conception of the contents of experiences, Caitlin’s experience is a *green*-experience (in our technical sense). On the other hand, we just saw that Caitlin’s belief (of the couch, that it is green) is not justifiably had. Hence, by repeating the familiar argument we again reach the conclusion that ER2 is false.

2.6 A LOOK AHEAD

If ER must be abandoned, what should we replace it with?

Let us begin by asking ourselves why Alice (at stage 4) and Bob (at stage 2) were able to justifiably believe, of the couch, that it is red. I submit that the simplest answer to this question is that they both had the ability to tell a red thing when they experienced it. In other words, their beliefs were justified because upon having the relevant experiences, they exercised their ability to tell (on the strength of these experiences) that the experienced thing is red. This simple answer fits nicely with the *Caitlin* scenario: Why was Caitlin (at stage 2) *unable* to justifiably believe, of the couch, that it is red? Precisely because Caitlin *lacked* any capacity to tell (on the strength of an experience) that the experienced thing is red.

It appears, then, that we have the makings of simple and intuitive answers to the two questions with which we opened this chapter: How can we form justifiable beliefs on the strength of experiences? In this way: Our experiences present us with (experienced) entities. If we have the requisite capacities to tell what these presented entities are, we can form justifiable beliefs on the strength of being experientially presented with them. And why is it that certain justifiable beliefs, but not others, can be formed on the strength of certain experiences, but not others? First, because different experiences present one with different entities. Second, because different persons have different capacities to tell what is the case on the strength of experiences.

This hastily sketched view has two advantages: First, unlike any version of ER, it delivers the correct results in all three of our scenarios. Second, it is simple, intuitive, and—because it does not appeal to the contents of experiences—free of commitments to a theory of content. A view along the lines of the one sketched here, I believe, should be the one we embrace.

In the next chapters I say more about the ultimate shape this view should take. I take hallucinations (but not illusions) to pose the most pressing problem before it. The problem is this: Hallucinatory experiences, it seems, present one with nothing. But if they present nothing, the hallucinator's capacities to tell what the presented entities are must remain idle. And on the present view it follows from this that it is impossible to form justifiable

beliefs on the strength of hallucinations. But it also seems quite possible to form justifiable beliefs on the strength of hallucinations. Contradiction. This problem is not a new one, and is shared by other accounts of a broadly relational spirit, such as Bill Brewer's (2011), John Campbell's (2002), Anil Gupta's (2006), Johannes Roessler's (2009) and Charles Travis' (2013a). I am unsatisfied with present answers to this problem, and I try to answer it along new lines.

3.0 HOW WE LEARN FROM PERCEPTION

This chapter develops a relationalist answer to the question, “how do we knowledgeably or justifiably form beliefs on the strength of perceptual experiences?” I hope the answer, which is inspired by the work of John L. Austin, will be readily accepted as both plausible and highly intuitive. Later chapters will reveal that the answer has a further virtue: It also tells us how we knowledgeably or justifiably form beliefs on the strength of *hallucinatory* experiences. But I am getting ahead of myself.

The view this chapter develops is, as I said, a relationalist one. But this does not mean that the view is inconsistent with the representationalist view, that experiences are evaluable for truth (veridicality, accuracy) at (centered) possible worlds. Indeed, I believe the two views are consistent. So the sense in which the view here developed is relationalist is only this: The view essentially relies on the presentational nature of experience, and is entirely silent on whether experiences are representational states. Beyond this, my objection to representationalism is restricted to my objection to Epistemic Representationalism (ER), which I presented in chapter 2.

The structure of this chapter is as follows: Section 3.1 presents how everyman talks and thinks about perception based acquisition of knowledge, exposing how we implicitly conceive of the matter. This implicit conception constitutes the heart of the epistemological view this chapter develops. Sections 3.2 and 3.3 then regiment the two key notions our implicit conception appeals to. Section 3.2 regiments the notion of a perceptual experience. Section 3.3 regiments the notion of a capacity to tell. Section 3.4 then draws on the work of the previous two sections and provides a general answer to the question “how can we knowledgeably or justifiably form beliefs on the strength of perceptual experiences?” Section 3.5 then highlights the appeals of the view which emerges by applying it to a particularly

difficult case. Section 3.6 concludes.

Before starting, two notes on terminology: (i) As the reader probably noticed, when I use the term “experience” without the modifier “perceptual” preceding it, I do so to collectively refer to both perceptual and non-perceptual (e.g., hallucinatory) experiences. This practice might now become confusing (especially in sections 3.3 and 3.4), since this chapter is only concerned with perceptual experiences. Nevertheless, the reader should keep in mind that unmodified occurrences of “experience” are not coextensive with “perceptual experience”. This terminological point will be relevant in chapter 5, where the ideas of this chapter will be generalized and where the fact that “experience” has a broad extension will be explicitly relied upon. (ii) Throughout this essay, I never assume that experiences have phenomenal characters. So, if zombies or blindsighters perceive things, then they enjoy (perceptual) experiences as well. It is important to keep this in mind, especially in section 3.3.

3.1 OUR IMPLICIT CONCEPTION

In his wonderful essay “Other Minds”, John L. Austin illustrates how we naturally answer questions of the form “how do you know that p ?” The illustration is illuminating, and deserves to be quoted at some length (1970, p. 79-80, statements are renumbered):

Suppose I have said “There’s a bittern at the bottom of the garden”, and you ask “How do you know?” [M]y reply may take very different forms:

- (a) I was brought up in the fens
- (b) I heard it
- (c) The keeper reported it
- (d) By its booming
- (e) From the booming noise
- (f) Because it is booming.

... That is, I may take you to have been asking:

- (1) How do I come to be in a position to know about bitterns?
- (2) How do I come to be in a position to say there’s a bittern here and now?
- (3) How do (can) I tell bitterns?
- (4) How do (can) I tell the thing here and now as a bittern?

The implication is that in order to know this is a bittern, I must have:

- (5) been trained in an environment where I could become familiar with bitterns
- (6) had a certain opportunity in the current case
- (7) learned to recognize or tell bitterns
- (8) succeeded in recognizing or telling this as a bittern.

Austin’s careful investigation of the presuppositions which guide the various ways in which we answer the question “how do you know” uncovers our implicit conception of non-inferential perception based acquisition of knowledge. According to this conception, for any subject *A* and proposition *p*, if *A* non-inferentially comes to know the *p* on the strength of perceptual experiences, *A* does so in the following way: First, prior to *A*’s coming to know, *A* undergoes some training. We allude to such training in making reports such as (a), and saying that one had such training can constitute answering concerns of the type raised in (1) in the way suggested by (5). Second, the upshot of the relevant training is that *A* acquires a *capacity* to tell things of a certain form.¹ We report that someone has such a capacity when we utter things like, “John can tell a bittern when he hears its booming”. This last utterance ascribes to John a capacity to tell things of the form *THAT* – *IS* – *A* – *BITTERN* on the grounds of hearing the booming of bitterns. Saying that one has such a capacity can constitute answering concerns of the type raised in (3) in the way suggested by (7). Third, once *A* has a capacity to tell certain things, perceptual experiences of the relevant kind constitute opportunities for *A* to *exercise* the capacity. These opportunities are what we report on in making reports such as (b), and what we allude to in making reports such as (d)-(f). Saying that one has such experiences can constitute answering concerns of the type raised in (2) in the way suggested by (6). Fourth, and finally, *A*’s non-inferentially coming to know the *p* on the strength of perceptual experiences is what a successful exercise of the relevant capacity culminates in, when *A* performs the exercise owing to those perceptual experiences. We allude to such exercises in saying things like, “John knows that that is a bittern because he heard its booming and thereby told it is a bittern”. Such reports can answer concerns of the type raised in (4) in the way suggested by (8).

Note three important features of our implicit conception of non-inferential perception based acquisition of knowledge: First, acquisition of such knowledge requires having capaci-

¹The first two points preclude having innate capacities to tell. But as the reader will observe, the view developed in the remainder of this essay is not committed to this consequence.

ties to tell. Second, the perceptual experiences which constitute opportunities for the exercise of these capacities are characterized in presentational terms. For example, the experience that provides John with the opportunity to know that that's a bittern is the experience we report on by saying that John hears the bittern's booming. John's hearing the booming of the bittern is an experience that presents John with a certain state of affairs (or, event). The presentation is accomplished by John's bearing the relation of hearing to the event. And the event itself is the particular booming of the particular bittern. Third, non-inferential perception-based acquisition of knowledge is the result of a successful exercise of the relevant capacity, owing to the appropriate presenting perceptual experiences. Putting the three points together, it emerges that our implicit conception is that A non-inferentially comes to know that p on the strength of perceptual experiences E just in case (i) A has perceptual experiences E , (ii) A has a capacity C to tell things of p 's form which is responsive to perceptual experiences of E 's kind, and (iii) A successfully exercises C owing to E .

So much for the description of our intuitive conception. What philosophical weight we should attach to this conception is a different matter. So let us ask: What is a philosopher to make of this implicit conception of ours? One option might be to dismiss it as hopelessly erroneous folk theory. This is not the option that I am about to recommend. I see no call for pessimism about this particular piece of folk theory. Rather, I believe we should try to provide a *mild rational reconstruction* of it. I am optimistic that such a reconstruction will be both highly intuitive and philosophically compelling. And since the conception is quite abstract and minimal in its empirical commitments, I believe that its reconstruction can be made to fit within a scientifically informed worldview.

In the rest of this chapter, our implicit conception will be made the cornerstone of a presentational account of how we knowledgeably or justifiably form beliefs on the strength of perceptual experiences. The account will attempt to reconstruct, clarify and extend this conception. And since the cornerstone conception is both intuitive and simple, I will call the account that will emerge the "Simple Picture" (or "SP", for short).

3.2 ON PERCEPTUAL EXPERIENCES

The first task before SP is to elucidate the thought that perceptual experiences are presentational mental states. In the introduction, that thought was rephrased as the thought that perceptual experiences are mental states that present worldly items to whoever has them. But this rephrasing is hardly an elucidation. To improve on it, let us—in an Austinian spirit—examine a certain way of using a certain class of transitive psychological verbs.

Consider a certain way of using the construction “*A* sees *x*”. This way of using the construction can be characterized by its having the following features:

1. On the intended use of the “*A* sees *x*” construction, what one says in uttering “*A* sees *x*” does not entail what can be expressed by uttering “*A* believes *x*”, “*A* is inclined to believe *x*”, “*A* understands *x*”, or “*A* knows *x*”. Compare: If I utter “Bob sees that John is wearing a red shirt”, what I thereby say *does* entail that Bob either knows, believes or recognizes that John is wearing a red shirt. This is not the way of using the “*A* sees *x*” construction that I have in mind. On the other hand, if I utter “Bob sees John’s wearing of a red shirt”, what I thereby say does *not* entail what can be expressed by uttering “Bob believes (or, is inclined to believe, ...) John’s wearing of a red shirt”. Such an utterance is ill formed. But more importantly, what I say in uttering “Bob sees John’s wearing of a red shirt” does not even entail that Bob believes (or, is inclined to believe, ...) *that* John is wearing a red shirt. In fact, what I say does not entail that Bob has *any* beliefs about John or about his shirt. Furthermore, what I say can be true even if Bob is completely unable to recognize John as John, or shirts as shirts.
2. On the intended use of the “*A* sees *x*” construction, what one says in saying “*A* sees *x*” cannot be true if *x* is something that can be thought. To illustrate, suppose you and I are playing chess. The game progresses, and you—being a far better player than I—are about to mate me. It is your turn to play, but before you move you say: “I have a puzzle for you. Where should I place my queen to mate you?” I think a little bit, and recognize that you should put your queen at C6 to mate me. You notice this and say “*now* you see the solution to the puzzle”. You thus used the “*A* sees *x*” construction in saying that I see that you should put your queen at C6. So the solution to the puzzle

is something thinkable, namely, that you should put your queen at C6. So your use of the construction “ A sees x ” is *not* the intended use. What sorts of things *can* be the references of “ x ” when one uses the “ A sees x ” construction in the intended way and in saying something true? Here is a very partial list: “ x ” can refer to scenes, locations, spaces, edges, material objects or persons. (All these can be called “objects”.) “ x ” can refer to colors, shades, lengths, shapes, sizes, motions, or facial expressions. (All these can be called “properties”.) “ x ” can refer to relative lengths, relative positions, relative sizes or relative distances. (All these can be called “relations”.) And finally, “ x ” can refer to certain states of affairs (or, events), i.e., to certain things’ having certain properties, or bearing certain relations to each other, at certain times and places. For example, one can see such things as your friend car’s being red, that bird’s current flying, the nearby pool’s being deep, or John’s wearing of a red shirt yesterday.

3. Intended uses of the “ A sees x ” construction are substitutionally transparent. This means that necessarily, if (i) using the construction in the intended way one utters “ A sees x ” at a context c , and (ii) x is identical to y , then were one to have used the construction in the the intended way in uttering “ A sees y ” (at c), one would have said something with the same truth value. For example, “Ovid sees Hesperus” entails “Ovid sees Phosphorus”. Similarly, “Ovid sees Hesperus’ shining this evening” entails “Ovid sees Phosphorus’ shining this evening”.
4. Intended uses of the “ A sees x ” construction are existentially transparent. This means that what one says in uttering “ A sees x ” (where one uses the construction in the intended way) entails that x exists or obtains. Here “exists” and “obtains” should be read tenselessly, i.e., as short for “exists, has existed or will exist” and “obtains has obtained or will obtain” respectively. For example, “ A sees a distant star” entails that A sees something that exists, has existed or will exist. (Note that in cases where the star is sufficiently far away, it may have ceased to exist by the time A sees it.) Similarly, “ A sees a distant star’s going supernova” entails that the star is going, has gone or will go supernova.²

²One might try to resist the claim that the “ A sees x ” construction is existentially transparent by noting that people sometimes say things like “I see a pink elephant”. But I do not claim that *all* ways of using “ A sees x ” are existentially transparent. I merely claim that there is *a* (well recognized) way of using “ A sees

I take these four observations to be uncontroversial claims about a certain way of using the “*A sees x*” construction which is already familiar to the reader, simply in virtue of her being at home with the verb “to see”. I am therefore about to rely on the reader’s familiarity with such uses of the construction in introducing the notion of a seeing: A *seeing* is a mental state that can be reported on by the outlined way of using the “*A sees x*” construction. In the rest of this essay, let us use the “*A sees x*” construction only in the seeings-reporting way. And if *A* sees *x*, let us call *x* the “*target*” of the seeing, or say that the seeing is a seeing “*of*” *x*.

Just as we introduced seeings by outlining four features of seeing-reporting uses of the “*A sees x*” construction, so can we introduce hearings, smellings, tastings and (tactile) feelings by outlining analogous features of the constructions that report on them. In other words, there are familiar ways of using the “*A hears x*”, “*A smells x*”, “*A tastes x*”, and “*A tactilely feels x*” constructions which (i) do not entail what can expressed by uttering “*A believes x*”, “*A is inclined to believe x*”, “*A understands x*”, or “*A knows x*” (ii) cannot be true when “*x*” refers to a thinkable, (iii) are substitutionally transparent and (iv) are existentially transparent. Hearings, smellings, tastings and (tactile) feelings are the mental states that can be reported on by these ways of using the constructions. In the rest of this essay, let us use the constructions only in the hearings-, smellings-, tastings-, and (tactile) feelings-reporting ways. And if *A* hears [smells, ...] *x*, let us call *x* the “*target*” of the hearing [smelling, ...], or say that the hearing [...] is a hearing [...] “*of*” *x*.

[What about the so-called “inner” senses, e.g., the senses of pain, hunger, balance, or body part location and movement? These senses can be treated analogously to the so-called “external” senses like sight. For example, we can introduce the notion of a *hurting* by appealing to the kind of mental state one reports on in uttering things like “*A feels the pain in her foot*”, or “*A feels her foot pain’s throbbing*”. Similarly, we can say that a *balancing* is the kind of mental state reported on in uttering “*A feels her falling down*”. We can even say that a *self-locating* is the kind of mental state reported on in uttering “*A feels her arm’s*

x” which it is existentially transparent. Those who assert things like “I see a pink elephant” are using the construction in a manner other than the intended one. I would, furthermore, point out that assertions like “I see a pink elephant” are to be regarded as non-standard, since it is appropriate to resist them by noting that one cannot strictly speaking see what is not physically before one’s eyes.

being behind her back”. Such locutions have features analogous to those of seeing-reporting locutions. Still, I will not attempt a full treatment of the “inner” senses here, as different ones do raise different special issues. For example, there would be need to consider the ramifications of the (purported) fact that there is no proprietary way it feels to *A* when *A* feels her arm’s being behind her back. And there would be need to consider the ramifications of the (purported) fact that one can report on a hurting not just by saying “*A* feels the pain in her foot”, but also by saying “*A*’s foot hurts”. Discussion of these issues will take us too far from our present concerns.

Once we have grasped what seeings, hearings etc. are, we can group all these mental states under the title “*perceivings*” or “*perceptual experiences*”. More explicitly, we say that:

(SP1) A perceiver, *A*, perceives *x* just in case *A* sees *x*, or smells *x*, or hears *x*, or ...

We then go on to say that a perceiving, or a perceptual experience, is a mental state that can be reported on by saying, of some *A* and *x*, that *A* perceives *x*. It follows that a perceiving is a mental state that consists in a perceiver’s perceiving a target - something which the perceiving is a perceiving of. This way of elucidating the notion of a perceptual experience allows us to give a simple and concrete meaning to the thought that perceptual experiences are presentational mental states: To say that *A* is perceptually presented with *x* is just to say that *A* perceives *x*.

3.3 ON CAPACITIES TO TELL

The second task before SP is to elucidate our implicit conception of capacities to tell things on the strength of experiences.

To begin, notice that a report like “John can tell a bittern when he hears its booming” assigns to John a capacity characterized in two ways. First, it is characterized by the sorts of things John can know by exercising the capacity. What John can thus come to know are things John can report on by uttering “that’s a bittern”. Second, the capacity is characterized by the sorts of experiences which constitute grounds for knowledge gainable by exercising

the capacity. These experiences are those reportable by uttering “John hears a bittern’s booming”.

To generalize from this example, we need a way of specifying, first, classes of things one can know, and second, kinds of experiences one can have.

The *second* of these tasks is not complicated. A kind of experience K can be specified thus: Start by specifying a type of experience by determining an experiential type (e.g., seeing, hearing, etc.) and a target or a type of target appropriate to that experiential type.³ We can thus speak of the types of experiences that constitute seeings of Barack Obama, hearings of bitterns or their boomings, or tactile feelings of something’s being rough. In the next step, let kind K be a collection of types of experience.

The *first* task is slightly more involved. Though it is not essential, it will be *much* easier for us to specify classes of things one can know if we allow ourselves to presuppose a standard view. According to this view: (i) To believe or to know something is to bear a certain belief- or knowledge-relation to a proposition. (ii) Propositions are structures (or logical forms) in which certain constituents (call them “atomic concepts”) occur. And (iii) the atomic concepts occurring in a proposition, along with the structural relations between them, determine the proposition’s truth conditions.⁴ Presupposing this standard view allows us to specify a class of things one can know by specifying a class of propositions. Such a class can in turn be specified by specifying a propositional form had by all and only propositions of that class. We shall say that two propositions have the same form iff the same atomic concepts occur in the same positions in the two propositions, except perhaps for atomic context-sensitive concepts, which may be different though they must still be of the same semantic kind.

Our way of specifying classes of things one can know can be illustrated by two examples: (i) I make judgments with different truth conditions by judging on Sunday and on Monday that the Sun is setting. Sunday’s judgment is true in virtue of the Sun’s setting on Sunday. Monday’s judgment is true in virtue of the Sun’s setting on Monday. On our way of classifying things one can know, one accounts for this fact by saying that although both judgments

³We will later add to the list of experiential types also visual hallucinations, auditory hallucinations, etc.

⁴I elsewhere expand on this standard view, articulating it in a way that reveals both that it is indeed a standard view, and that it is largely independent of one’s views about the *content* of mental states.

consist in my coming to bear the belief- or knowledge-relation to propositions of the form *THE SUN – I – SETTING*, the two judgments involve distinct propositions, which can be symbolized as the propositions *THE SUN – IS₁ – SETTING* and *THE SUN – IS₂ – SETTING*. The tensed concept *IS₁* is involved in Sunday’s judgment and makes it the case that the truth of this judgment depends on what occurred on Sunday. The tensed concept *IS₂* is involved in Monday’s judgment and makes it the case that the truth of this judgment depends on what occurred on Monday. (ii) I make judgments with different truth conditions by judging of distinct bitterns - Polly and Birdy - that they are bitterns. On our way of classifying things one can know, one accounts for this fact by saying that although both judgments consist in my coming to bear the belief- or knowledge-relation to propositions of the form *I – I – A – BITTERN*, the two propositions involved have distinct concepts—*THAT₁* and *THAT₂*—occurring in the demonstrative *I* position. (For simplicity I ignore the tensed *I* position.) The demonstrative concept occurring in the proposition involved in one judgment makes it the case that its truth depends on the kind of thing Polly is. The distinct demonstrative concept occurring in the proposition involved in the other judgment makes it the case that its truth depends on the kind of thing Birdy is.

[Note that the proposition *THAT – IS – A – BITTERN* which concerns Birdy is *not* of the same form as the proposition *HERE – IS – A – BITTERN* which concerns Birdy’s location. It is true that, with the exception of the context-sensitive concepts, the same concepts occur in the same positions in the two propositions. Nonetheless, the propositions have distinct forms because the context-sensitive concepts *HERE* and *THAT* have distinct semantic kinds. I leave to another essay the task of giving a fuller treatment of the question of when two context sensitive concepts belong to the same semantic kind.]

Let us return to topic. We saw that our conception of a capacity to experientially tell is that of a capacity which can be characterized by the sorts of things it affords knowledge of, and by the kinds of experiential grounds owing to which it can afford knowledge. We can now better express this idea by saying that a capacity to experientially tell is a capacity which can be characterized (i) by the class of propositions one can come to know by successfully exercising it, and (ii) by the class of experiences which enable such successful exercises.⁵

⁵I believe that characterizing capacities to tell in the proposed way will do the most to help us understand

It remains to say how particular successful exercises of a capacity to tell, owing to particular experiences of the relevant kind, lead to knowledge of particular propositions of the relevant form. Saying this, however, requires drawing a distinction between capacities to tell which afford knowledge of propositions in which no experiential demonstrative concepts occur (e.g., the form *THE SUN – I S – SETTING*), and capacities to tell which afford knowledge of propositions in which a single atomic experiential demonstrative concept occurs (e.g., the form *I H A T – IS – A – BITTERN*). [Capacities to tell which afford knowledge of propositions in which two or more atomic experiential demonstrative concepts occur can be treated in analogous fashion to capacities of the latter variety and so will be left alone in the remainder of this chapter.]

To begin, consider capacities to tell of the first variety:

(SP2a) If no atomic experiential demonstrative concept occurs in a proposition of the form \mathcal{P} , then $C(\mathcal{P}, K)$ is a capacity to tell iff

- (i) the capacity can be successfully exercised,
- (ii) any successful exercise of the capacity by a subject A culminates in A 's non-inferentially coming to know some proposition p of the form \mathcal{P} owing to some (on-line or remembered) experiences E of the kind K , and
- (iii) any exercise of the capacity by a subject A culminates in A 's non-inferentially coming to believe some proposition p determined by (a) its having the form \mathcal{P} and (b) the circumstances under which A performs the exercise.⁶

This characterization uses some quasi-technical terminology, which will be explained more fully soon. In the meantime, two examples will help us find our bearings. Consider a capacity to experientially tell $C(\textit{THE SUN – H A S – SET – T O D A Y}, \textit{seeings of a night sky's being dark})$. This is a capacity that enables a subject to know that the sun has set on various days owing to seeings whose targets are those events which consist of a night sky's being dark. If I successfully exercise this capacity on Sunday night, owing to my on-line seeing of the night sky's being dark, I end up knowing the proposition *THE SUN – HAS – SET – TODAY*, where the concept *HAS* refers to the past relative to Sunday night, and the concept *TODAY*

why a particular successful exercise of such a capacity culminates in the knowledge of one thing rather than another. But this does not mean that there are no additional ways to characterize and individuate such capacities. In fact, for all I have said, all of the capacities to tell that I discuss may be subsumed under a further “general capacity to tell things on the strength of experiences”. Whether this is so or not is, however, irrelevant for my purposes here.

⁶Capacities to tell that can relate one to a single proposition—one in which no context sensitive concepts at all occur—constitute (if they exist) a degenerate case of capacities of this variety.

refers to Sunday. These concepts have these references (i) because of their respective semantic types, and (ii) because I exercised my capacity to tell on Sunday night.

Now consider a different case. Suppose on Friday night I confusedly take a (remembered) experience I had on Tuesday to be an experience I had earlier that night. The experience in question is a seeing of the night sky's being dark on Tuesday. Suppose further that in virtue of this confusion I do something I shouldn't: On Friday night I exercise the aforementioned capacity in a way which culminates in my believing the proposition *THE SUN – HAS – SET – TODAY* owing to the (remembered) seeing of Tuesday's event. [Since I undertake this exercise because I have a (mistaken) belief about when I had the relevant experience, the resulting belief is formed not just owing to the experience, but also to the mistaken belief.⁷ And since my mistaken belief is involved, my exercise is unsuccessful and culminates not in knowledge, but in mere belief.]

The thing to note is that the proposition *THE SUN – HAS – SET – TODAY* is true iff the sun has set on *Friday*. The belief's truth conditions are therefore not sensitive to Tuesday's sunset, despite the fact that the seeing which grounds my belief *is* a seeing of an event which occurred on Tuesday. This shows that the references of the concepts *HAS* and *TODAY* is determined by the circumstances under which I exercised my capacity to experientially tell, and not the circumstances in which I had the relevant experience.

Consider next capacities to tell of the second variety:

(SP2b) If a single atomic experiential demonstrative concept occurs in a proposition of

⁷This is a good opportunity to flag that exercises of capacities to tell can culminate in judgments made owing not *just* to certain experiences, but *also* to certain beliefs. One way for this to happen is when the capacity to tell is such that one has it in virtue of having certain knowledge or beliefs. Exercises of such a capacity culminate in judgments made owing not *just* to the relevant experiences, but also to the enabling knowledge or beliefs. For example, consider my capacity to tell my car is doing 60 mph upon seeing its speedometer's so indicating. I only have this capacity in virtue of my knowing that when my car's speedometer indicates 60, my car is doing 60. Therefore, when I successfully exercise the capacity, I come to know that my car is doing 60 owing both to my seeing the speedometer's so indicating and to my knowing that when my car's speedometer indicates 60, my car is doing 60.

Note: I do not *infer* that my car is doing 60 from my beliefs about what the speedometer indicates and about what its indications suggest. I judge non-inferentially. It is only that my ability to do that would be lost if I ceased to believe that when my car's speedometer indicates 60, my car is doing 60.

There are other ways for an exercise of a capacity to tell to culminate in a judgment made owing not just to experiences, but also to beliefs. These ways have to do with features of the particular exercise involved, such as whether the exerciser undertook the exercise in virtue of having certain beliefs. If so, the judgment the exercise culminates in will be made owing to those beliefs as well. The example in the main text illustrates just this point.

the form \mathcal{P} , then $C(\mathcal{P}, K)$ is a capacity to tell iff

- (i) the capacity can be successfully exercised,
- (ii) any successful exercise of the capacity by a subject A culminates in A 's non-inferentially coming to know some proposition p of the form \mathcal{P} owing to some (on-line or remembered) experiences E of the kind K , and
- (iii) any exercise of the capacity by a subject A culminates in A 's non-inferentially coming to believe some proposition p determined by (a) its having the form \mathcal{P} , (b) the circumstances under which A performs the exercise, and (c) the fact that an experiential tracking effort F controls the atomic experiential demonstrative concept D taking the demonstrative position in \mathcal{P} .

This characterization uses even more jargon than the last one. But before we turn to that, an example. Consider a capacity to tell $C(\mathcal{T H A T} - \mathcal{I S} - A - \mathcal{B I T T E R N}, \text{hearings of a bittern's booming})$. This is a capacity that enables a subject to know, of various bitterns, that they are bitterns, owing to hearings whose targets are those events which consist of some bittern's booming. Exercises of this capacity involve efforts to track a unique experienced target. If the exercise of the capacity is successful, there is a unique target actually tracked, and it serves as the reference of the demonstrative concept (assuming all else is in order). Thus, if I successfully exercise the capacity at t owing to my on-line hearing of Polly's booming (which involves me in an experiential tracking of Polly), I end up knowing the proposition $\mathcal{T H A T} - \mathcal{I S} - A - \mathcal{B I T T E R N}$, where the concept $\mathcal{T H A T}$ refers to Polly, and the tensed concept $\mathcal{I S}$ refers to some stretch of time which includes t . The concept $\mathcal{T H A T}$ has its reference (i) because of its semantic type, and (ii) because the tracking effort that controlled it was a tracking of Polly. The concept $\mathcal{I S}$ has its reference (i) because of its semantic type, and (ii) because I exercised the capacity to tell at t .

Now consider a different case. Suppose I successfully exercise the capacity at t owing to my *remembered* hearing of Polly's booming at t' . Thanks to this successful exercise, I end up knowing the proposition $\mathcal{T H A T} - \mathcal{W A S} - A - \mathcal{B I T T E R N}$, where the concept $\mathcal{T H A T}$ refers to Polly, and the tensed concept $\mathcal{W A S}$ refers to the past relative to t . The concept $\mathcal{T H A T}$ has its reference (i) because of its semantic type, and (ii) because the tracking that controlled it was a tracking of Polly. The concept $\mathcal{W A S}$ has its reference (i) because of its semantic type (where the past-tense inflection reflects that I undertook the exercise while recognizing that the relevant experience is a remembered one), and (ii) because I exercised

the capacity to tell at t .

Our examples reveal a distinction between two categories of context sensitive concepts: experiential demonstrative concepts (i.e., concepts controlled by experiential tracking efforts) and other context-sensitive concepts. The references of experiential demonstrative concepts are fixed by their semantic kinds and by facts about what, if anything, one tracked by the controlling experiential tracking effort. Such tracking efforts constitute experience-based demonstrations to which experiential demonstrative concepts can be indebted. The category of other context-sensitive concepts includes concepts that are not similarly indebted to an experience-based demonstration. Their references are rather fixed by their semantic kind and by facts about the circumstances under which the subject first tokened them.

(SP2a) and (SP2b) equip us with a fuller characterization of capacities to tell, but only by using a number of terms (“owing to”, “inferentially coming to know”, “a successful experiential tracking effort”, “control” and “successful exercise”) which deserve further examination. So let us examine them.

3.3.1 “Owing To”

The expressions “ A believes (knows) that p owing to experiences E ” and the expression “ A (knowledgeably) judges that p on the strength of experiences E ” are used interchangeably in this essay. I elucidate them by noting that A believes (knows) that p owing to experiences E iff a correct (even if partial) answer to the question, “why does A (knowledgeably) judge that p ?” can be given by saying that A does so because A had experiences E [where “because” takes its specific epistemic meaning, rather than a generic causal meaning, or some other non-epistemic meaning].

3.3.2 “Inferentially Coming to Know”

In the context of this chapter, we shall say that A *inferentially* comes to believe (know) that p iff A comes to believe (know) that p by *consciously* inferring that p from other (knowledgeable) beliefs. We shall further say that A *non-inferentially* comes to believe (know) that p iff A comes to believe (know) that p but does not do so inferentially. Here

“inferring” need not be logical inferring. It may also be probabilistic, inductive, etc.

3.3.3 “Experiential Tracking Effort”, “Success” and “Control”

We all know what it is to experientially track something. We also all know that we can fail to track something without noticing it. How can we capture what is going on with us in these two cases? Here is one suggestion: Let us say that A engages in *an experiential tracking effort* F iff F consists of its being the case that (i) at some instant or duration of time t , A has a collection of experiences $e_1, e_2, \dots, e_i, \dots$ with respective targets $x_1, x_2, \dots, x_i, \dots$, and (ii) for any i , and not owing to any beliefs A has or to any reasoning A engages in, it is to A as if there is some unique thing x such that $x = x_i$.⁸ Let us also say that the experiences $e_1, e_2, \dots, e_i, \dots$ involved in an experiential tracking effort F are the experiences that “guide” F . With these two definitions in hand, we can say that F is a *successful* experiential tracking effort iff there is a unique thing x which is the target of all the experiences that guide F . Finally, we can say that A experientially tracked x iff A engaged in some successful experiential tracking effort F , and x was the unique target of all the experiences that guided F . A few examples will make this clearer.

(1) Suppose that I both hear and see my favorite bittern, Polly, and that it is to me as if I see and hear a single thing. In that case I have successfully tracked Polly, and I was guided by my sighting and hearing of her.

(2) Suppose I am in the fens, looking at a grassy marsh. In the marsh are two birds - a bittern and a heron. The heron - Tweety - is visible to me and silent. The bittern - Polly - is booming but completely occluded by the grass. So I hear Polly (and Polly’s booming), and I see Tweety. Not realizing that there are two birds around, it is to me as if I see and hear a single thing.⁹ In this case I have not tracked anything, and my tracking effort is unsuccessful. It was, after all, guided by experiences of distinct birds.

(3) Suppose I see two bricks - a red one and a green one - flying past each other. I see

⁸I take no stand on how experiences are to be individuated, except for pointing out that experiences with distinct targets are themselves distinct. The definitions set out in this paragraph are designed to allow for this.

⁹It makes it easier to understand how it can be that way to me if we suppose that I have no capacity to tell bitterns, or any other kind of bird, by *sight*.

the red brick both before and after it passes the green one. It is to me as if I see a single red thing in both sightings. In this case I have successfully tracked the red brick, and I was guided by my two sightings of it.

(4) Suppose I see two bricks - a red one and a green one - flying towards each other. As they hit they exchange colors - the red one turns green and the green one turns red. I can't hear the impact because I have wax in my ears. Immediately after the bricks hit, they oddly start flying back away from each other. I see this too. Supposing all this happens against the appropriate background, the whole event looks to me just the way it looks to me when two bricks - a red one and a green one - fly past each other.¹⁰ So it is to me as if I see as single constantly colored red brick once to the left of another, and then to its right. In fact, one of the relevant sightings is of a red brick later to turn green, and the other sighting is of a different red brick, which was green earlier. In this case I have not tracked anything, and my tracking effort is unsuccessful. The effort was guided by experiences of distinct bricks.

It remains to explain the idea of *control*. Let us say that an experiential tracking effort *F* controls an atomic experiential demonstrative concept *D* iff:

- either (i) *D* refers to *x* because (a) *A* experientially tracked *x* by engaging in *F*, and (b) *x* is a possible reference for concepts of *D*'s semantic kind,¹¹
- or (ii) *D* refers to nothing because (a) *F* is unsuccessful, or (b) *A* experientially tracked something by engaging in *F*, but the thing is not a possible reference for concepts of *D*'s semantic kind.

This definition of control embeds in it the outline of a small theory of reference fixing for atomic experiential demonstrative concepts. I should like to say something to motivate this theory: An experiential demonstrative concept *D* refers (modulo considerations regarding *D*'s semantic kind) to what its possessor demonstrates in using *D*. But how should we de-

¹⁰To experience this yourself, see Arthur Shapiro & Gideon Caplovitz's demonstration #1, in their illusion *The Exchange of Features, Textures and Faces*, available at <http://illusionoftheyear.com/2011/the-exchange-of-features-textures-and-faces/>

¹¹Why is clause (b) required? Here is an illustration: Suppose that though I experientially track Polly by engaging in effort *F*, I mistake her for a man. Suppose I then think *HE - IS - AN - ODD - LOOKING - CHAP*, where *HE* is controlled by *F*. It is clear that in such a case *HE* refers to nothing. Clause (ii) explains why - a bird is not a possible reference for a concept of *HE*'s semantic kind. [In this context I should point out that there are ways of using the words "that" and "he" (as well as others) in ways that do not express experiential demonstrative concepts. Thus, if you say, "Bush was a wonderful leader", and I reply either with "that was no leader at all", or with "he was no leader at all", then my uses of "that" or "he" are anaphorically related to the subject of your assertion. They therefore express the concept *BUSH* rather than an experiential demonstrative one.]

termine what, if anything, she demonstrates? Given that D is an experiential demonstrative concept, a natural answer is to say that she demonstrates that which is the target of the experiences she associates with D . And she associates an experience with D just in case the experience guides the tracking effort to which her uses of D are indebted. If we call the tracking effort to which her uses of D are indebted the tracking effort that “controls” D , our line of thought leads to the characterization of control we gave above. Or it nearly leads to it. It remains to explain why D does not refer when the tracking effort it is indebted to is unsuccessful. The explanation is best given by example.

Consider again the example of Polly and Tweety. In that example, an (unsuccessful) tracking effort F is guided by a hearing of Polly the bittern and a sighting of Tweety the heron. Now suppose that I come to (mock) believe the propositions $THAT - IS - A - BITTERN$ and $I - SEE - THAT$, in which occurs a single concept $THAT$ controlled by F .^{12,13} Since both occurrences of $THAT$ are indebted to a single tracking effort, guided by a single set of experiences, they should have the same reference, if any. But what is this reference? If the reference is Tweety, then $THAT - IS - A - BITTERN$ is false, while $I - SEE - THAT$ is true. If the reference is Polly, then $THAT - IS - A - BITTERN$ is true, while $I - SEE - THAT$ is false. Neither interpretation seems preferable to the other. So it is best to say that $THAT$ does not refer at all.

Another consideration leading to the same conclusion is this: Suppose I continue tracking effort F well into a time in which I both continue to hear Tweety and see Polly as she is flying far away. I then realize that there were not one but two birds around. Realizing this, what should I do with the propositions $THAT - IS - A - BITTERN$ and $I - SEE - THAT$? Well, it seems that my internal monologue should go something like this: “There was no single bird of which I earlier thought that it is a bittern. There was also no single bird of which I earlier thought that I see it. I confused two birds for one in thinking these thoughts.

¹²Reminder: I have no capacity to tell bitterns, or any other kind of bird, by *sight*.

¹³If a proposition has no truth conditions, some thinkers prefer to say that one cannot bear the belief-relation to it. Their thought is that if A bears the belief-relation to a proposition, A thereby comes to believe something; and that since the proposition has no truth conditions, A can believe nothing by being related to it. Such thinkers do admit, however, that when a proposition has no truth conditions, things can be *as if* A stands in the belief-relation to it. On their preferred way of talking, A mock believes that p iff p has no truth conditions, and it is as if A stands in the belief-relation to p . In this chapter I try to accommodate this way of thinking.

So both these thoughts were likewise confused. Good thing I no longer think them.” This, I think, is the way I should be talking with myself. And the internal monologue suggests the following semantic analysis: First, since I should abandon both thoughts, neither proposition is true. Second, since both propositions are confused rather than mistaken, neither is simply false. And the best account of this is the neither proposition has definite truth conditions, as the concept *THAT* occurring in them does not refer.

Objection. Let’s rethink the Tweety and Polly example. You claim your concept *THAT* refers to nothing, because *F* is guided by experiences of both Polly and Tweety. But there is no call for that conclusion. Suppose that, before you realized your confusion, you expressed your thoughts by uttering “that is a bittern, and I see it”. If I then asked you what you meant by “that”, you would point to Tweety. This suggests two things. First, the fact that you would point to Tweety in response to my question indicates that you said something about Tweety. Second, given that you did say something about Tweety in expressing your thoughts, it follows that your thoughts too were about Tweety. Hence *THAT* refers to Tweety.

Reply. Both halves of the objection should be resisted. To begin, do I really say something about Tweety if, pointing at Tweety, I utter “that is a bittern, and I see it”? How does my uttering the word “that” while pointing at Tweety constitute my saying something about Tweety, rather than about the mereological sum of Tweety *and* Polly, or about nothing at all? A possible answer is that I say something about Tweety because, given my pointing gesture, a normal, rational addressee would take me to be saying something about Tweety. But this answer cannot be right. That an appropriate addressee would take me to be saying something about Tweety leaves entirely open the possibility that the addressee would be *mistaken* in so taking me. And indeed, the addressee would be mistaken. If the addressee is to correctly understand what I say, she must determine what I am talking about. And she cannot determine this from considering the direction I am pointing to. It is a familiar Wittgensteinian point that there are simply too many things in that direction, including the marsh, Tweety, Tweety’s beak, the mereological sum of Tweety and Polly, etc. So to determine what I am talking about, the addressee has to consider not just the orientation of my finger, but also what unique thing I have in mind. And since, *ex hypothesi*, I have failed to track a unique entity, I have no unique thing in mind. So the addressee should determine

that I am not talking about anything. If she determines otherwise, she is simply mistaken.

I turn to the second half of the objection. For the sake of argument, I now grant that I say something about Tweety in uttering “that is a bittern, and I see it”. I still do not see why anyone should think this shows that my concept *THAT* refers to Tweety. Plainly, what I say is one thing, and what I think is quite another, as I can *fail* to express my thoughts in speaking. A possible rejoinder may be this: “You grant that you are saying something about Tweety. But you are only succeeding in saying something about Tweety because you intend your use of the demonstrative term to refer to Tweety. And that you intend your use of the term to refer to Tweety indicates that your concept *THAT* refers to Tweety.” But the rejoinder is confused. If I am to intend my use of “that” to refer to Tweety, I must have some concept in my possession that both refers to Tweety and that I can rationally employ in intending, of Tweety, to refer to her by “that”. (To illustrate, I must have some concept *X* which refers to Tweety and which could occur in my rationally bearing the intending-relation to the action-content *TO-REFER-TO-X-BY-UTTERING- \lceil that \rceil* .) But, assuming that I have not recognized Tweety by subsuming it under any singularly referring attributive concept, it is difficult to see what that concept might be unless it is a demonstrative concept. And what is now under debate is precisely the question of whether I have in my possession a demonstrative concept referring to Tweety. So the rejoinder is question-begging. The rejoinder is furthermore mistaken, as previous arguments showed that I do not have such a demonstrative concept.

Objection [due to Millar (2008, p. 337-339)]. You have just given an account of experiential tracking efforts and of experiential demonstrative reference which makes essential use of the idea that experiences are mental states that have targets. But we can also account for the phenomena of experiential tracking and demonstrative reference without invoking such an explicitly relational conception of experiences. We can just say this instead: Experiences are non-relational mental states fully described by their phenomenal character. Still, it can be that an entity controls the phenomenal character of an experience of A's and installs in A potentialities for sub-doxastic discriminatory behaviors directed towards it. An entity's so doing constitutes its being picked out by A. The relation of picking out is a sensory (e.g., visual) relation. The phenomenon of experiential demonstrative reference (and perhaps that

of experiential tracking as well) should be accounted for in terms of this relation, rather than in terms of some “experiential” relation.

Reply. The objection suggests that there is some relation—called “picking out”—such that (i) it is a sensory (e.g., visual) relation, and (ii) for it to obtain between A and x is for x to (a) “control the phenomenal character of an experience” of A ’s and to (b) “install in A potentialities for sub-doxastic discriminatory behaviors directed towards” x . Given this characterization of the relation, one might ask oneself if there are any necessary conditions on A ’s bearing the relation to x . The answer is clear: A can only pick out x if A experiences x . If A does not experience (e.g., see) x , what could be meant by claiming that A bears a sensory relation to x ? Similarly, if A does not experience x , in what sense can it be x (rather than some other related object, property, relation or event) that controls the phenomenal character of an experience of A ’s? And finally, if A does not experience x , how could A ’s behaviors be directed towards x (rather than some other related object, property, relation or event)? It must be conceded, then, that A can pick out x only if A experiences x . So we are now in the following predicament: By the objection’s own lights, an account of the phenomenon of experiential demonstrative reference is to be given in terms of the picking-out relation, which in turn, can only obtain between two relata in virtue of the obtaining of an experiential relation between them. In other words, by the objection’s own lights, A can refer to x using an atomic experiential demonstrative concept only if A experiences x . This is precisely what the objection hoped to deny.

A possible rejoinder might concede that A can refer to x using an atomic experiential demonstrative concept only if A experiences x . The rejoinder might concede this, only to go on insisting that experiences are non-relational mental states fully described by their phenomenal character. The coherence of this insistence can be secured by claiming that the locution “ A experiences x ” does not report on the mental states which experiences are.

If this is all that the rejoinder means to insist on, the rejoinder is no threat to SP. The word “experience” can be used as a name for all sorts of mental states. SP uses it as a name for a type of presentational mental state, on which one can report by certain locutions which satisfy certain conditions (as described in section 3.2). The rejoinder insists on using the same word for another kind of mental state. If this is all the rejoinder means to insist on,

then its disagreement with SP reduces to a merely verbal one.

3.3.4 “Successful Exercise”

The notion of a *successful* exercise of a capacity to tell is a particularly important one: According to SP, a successful exercise of a capacity to tell always culminates in the non-inferential experience-based acquisition of *knowledge*.

An exercise of a capacity to tell is *successful* iff it is (i) not *reckless* and (ii) performed under *epistemically favorable circumstances*. Let us look at these two subsidiary notions in turn.

An exercise of a capacity to tell is *reckless* iff its undertaking demonstrates insufficient care for the veracity of the judgment that the exercise culminates in. One way to be reckless in exercising a capacity to tell is to exercise it under conditions the exerciser should appreciate are inhospitable. For example, you are being reckless if you exercise your capacity $C(\mathcal{T}\mathcal{H}\mathcal{A}\mathcal{T} - \mathcal{I}\mathcal{I} - A - \text{BITTERN}, \text{hearings of a bittern's booming})$ when the booming you hear is appreciably not quite distinctive of bitterns. Another way to be reckless in exercising a capacity to tell is to exercise it when one is alive to some reason for which the judgment the exercise culminates in should not be made. For example, you are being reckless if, knowing that a solar eclipse is supposed to take place sometime today, you judge that the sun has set on the strength of seeing the night sky's being dark. Here you are being reckless because you make your judgment while you have a good reason to suspect that the sun has not set, and that you are merely seeing the *day* sky's being dark as a result of an eclipse.

An exercise of a capacity to tell is performed under *epistemically favorable circumstances* iff it is performed under circumstances in which the judgment that the exercise culminates in is non-accidentally true. Here it is important to note that there are various kinds of epistemically *unfavorable* circumstances, which contribute in different ways to the production of false, merely accidentally true, or truth-valueless (or mock) beliefs. I will illustrate just three kinds of epistemically unfavorable circumstances.

First, there are circumstances that are epistemically unfavorable because they manifest

the *fallibility* of our capacities to tell. We say that a capacity to tell $C(\mathcal{P}, K)$ is *fallible* iff it can be exercised when A 's experiences are *not* of kind K . Many, if not all, of our capacities to tell are fallible in this sense. For example, my capacity $C(\mathcal{T}\mathcal{H}\mathcal{A}\mathcal{T} - \mathcal{I}\mathcal{S} - A - \text{BITTERN}, \text{hearings of a bittern's booming})$ is fallible because I can exercise it not only when I hear a bittern's booming, but also when I hear the a mere ringer for that booming. Similarly, my capacity $C(\text{THE SUN} - \mathcal{H}\mathcal{A}\mathcal{S} - \text{SET} - \mathcal{T}\mathcal{O}\mathcal{D}\mathcal{A}\mathcal{Y}, \text{seeings of a night sky's being dark})$ is fallible because I can exercise it not only when I see the night sky's being dark, but also when I see a ringer for that, like a window painted over to look like the dark night's sky.

If $C(\mathcal{P}, K)$ is fallible, and A 's exercise of it culminates in A 's judging that p owing to experiences E which are *not* of kind K , then A 's exercise occurred under epistemically unfavorable circumstances, and A 's belief is truth-valueless, false or merely accidentally true. For example, say an exercise of my capacity $C(\text{THE SUN} - \mathcal{H}\mathcal{A}\mathcal{S} - \text{SET} - \mathcal{T}\mathcal{O}\mathcal{D}\mathcal{A}\mathcal{Y}, \text{seeings of a night sky's being dark})$ culminates in my judging that the sun has set today owing to my seeing a window painted over to look like the dark night's sky. Say, furthermore, that this exercise occurs after sunset. In that case, my judgment is merely accidentally true.

Second, there are circumstances that are epistemically unfavorable because they lead one to *mishandle* the exercise of a capacity to tell $C(\mathcal{P}, K)$, even if the exercise occurs when one has experiences of kind K . I will illustrate this idea using two cases, both of which will be familiar from earlier in this section:

(1) Suppose on Friday night I confusedly take my (remembered) seeing of the night sky's being dark on Tuesday to be a seeing I had earlier that night. As a result, on Friday night I undertake a certain exercise of my capacity $C(\text{THE SUN} - \mathcal{H}\mathcal{A}\mathcal{S} - \text{SET} - \mathcal{T}\mathcal{O}\mathcal{D}\mathcal{A}\mathcal{Y}, \text{seeings of a night sky's being dark})$. This exercise culminates in my believing that the sun has set today (i.e., on Friday) owing to (i) the (remembered) seeing of Tuesday's event, and (ii) the (mistaken) belief that the recollected experience is one I had earlier that night.¹⁴ It is clear that this belief is merely accidentally true and not knowledgeable. The belief is accidentally true because it emerges from my mishandled exercise of my capacity. I mishandle this exercise because I undertake it in virtue of a false belief about the time at which I had the relevant experience. My having this false belief constitutes an epistemically unfavorable

¹⁴For more on this point, see footnote 7.

circumstance for that particular exercise.

(2) Suppose an (unsuccessful) tracking effort F of mine is guided by a hearing of Polly the bittern and a sighting of Tweety the heron. Suppose further that as I see Tweety and hear Polly I exercise my capacity $C(\mathcal{T}\mathcal{H}\mathcal{A}\mathcal{T} - \mathcal{I}\mathcal{S} - A - \text{BITTERN}, \text{hearings of a bittern's booming})$ in a way that involves F . This means that the exercise culminates in my coming to (mock) believe the proposition $\text{THAT} - \text{IS} - A - \text{BITTERN}$, where the concept THAT is controlled by F .¹⁵ It is clear that the (mock) belief is not non-accidentally true because it emerges from my mishandled exercise of my capacity. I mishandle this exercise because the exercise involves a tracking effort in which I fail to track a unique entity. My failure constitutes an epistemically unfavorable circumstance for the exercise.

Third, there are circumstances that are epistemically unfavorable because they deprive the relevant capacity of the status of a capacity to tell. Consider again my capacity $C(\mathcal{T}\mathcal{H}\mathcal{A}\mathcal{T} - \mathcal{I}\mathcal{S} - A - \text{BITTERN}, \text{hearings of a bittern's booming})$. If I move away from the fens to live in bittern-ringer county (where there are very few bitterns but lots of mechanical robots that look and boom just like bitterns), this capacity will lose the status of a capacity to tell. Instead, it will be demoted to a capacity to come to bear the belief-relation to propositions of the form $\mathcal{T}\mathcal{H}\mathcal{A}\mathcal{T} - \mathcal{I}\mathcal{S} - A - \text{BITTERN}$ upon hearing bittern boomings. Any exercise of this capacity occurring in bittern-ringer county will thus count as unsuccessful, and as one performed under epistemically unfavorable circumstances. To illustrate this, suppose that while living in bittern-ringer county I happen to hear the booming of one of the few actual bitterns there. Suppose further that, being unaware that I am surrounded by bitter-ringers, I non-recklessly exercise my belief-forming capacity. Through this exercise I come to believe, of the bittern, that it is a bittern. In this case, the fallibility of my capacity did not manifest itself, because I did hear a true bittern's booming. Still, the exercise of my capacity occurred under epistemically unfavorable circumstances, and culminated in an *accidentally* true belief. This accounts for the fact that I did not end up knowing, of the bittern, that it is a bittern.

To recapitulate: An exercise of a capacity to tell is *successful* iff it is (i) not reckless and (ii) performed under epistemically favorable circumstances. An exercise is not reckless iff its

¹⁵Reminder: I have no capacity to tell bitterns, or any other kind of bird, by *sight*.

undertaking demonstrates sufficient care for the veracity of the judgment that the exercise culminates in. An exercise is performed under epistemically favorable circumstances iff it is performed under circumstances in which the judgment that the exercise culminates in is non-accidentally true. This concludes our presentation of capacities to tell.

3.4 ON KNOWING, JUSTIFIABLY BELIEVING, AND RATIONALLY BELIEVING

We are now ready to examine SP's account of learning from experience:

(SP3a) *A non-inferentially comes to know that p owing to (on-line or remembered) experiences E iff for some capacity to tell $C(\mathcal{P}, K)$, A 's successful exercise of $C(\mathcal{P}, K)$ culminates in A 's coming to believe that p owing to E .*

(SP3b) *A non-inferentially comes to justifiably believe that p owing to (on-line or remembered) experiences E iff for some capacity to tell $C(\mathcal{P}, K)$, A 's non-reckless exercise of $C(\mathcal{P}, K)$ culminates in A 's coming to believe that p owing to E .*

(SP3c) *A non-inferentially comes to rationally believe that p owing to (on-line or remembered) experiences E iff A non-recklessly and without inference comes to believe that p owing to E .*

The reader will recall that earlier in this chapter we saw that our implicit conception of learning from perceptual experience suggests that A non-inferentially comes to know the p on the strength of perceptual experiences E just in case (i) A has perceptual experiences E , (ii) A has a capacity C to experientially tell things of p 's form which is responsive to perceptual experiences of E 's kind, and (iii) A successfully exercises C owing to E . We can now appreciate how (SP3a) captures this idea. It does so because it follows from (SP2a) and (SP2b) that if A successfully exercises a capacity to tell $C(\mathcal{P}, K)$, the resulting known proposition will have the form \mathcal{P} , and A will come to know it owing to experiences of kind K .

Theses (SP3b) and (SP3c) go beyond our implicit conception. Through these theses SP recognizes two epistemically meritorious statuses that a belief can have even if it falls short of knowledge. (SP3b) suggests that an experience-based judgment is justifiably and non-inferentially made just in case it emerges from a non-reckless exercise of a capacity to tell.

(SP3c) suggests that an experience-based judgment is rationally and non-inferentially made just in case it is made non-recklessly and without inference. If a non-inferential and non-knowledgeable experience-based judgment has either one of these epistemically meritorious statuses, then it was made under epistemically unfavorable circumstances.

There is a crucial difference between justifiably and merely rationally believing a proposition. If *A* merely rationally believes a proposition (non inferentially and owing to experiences), then although *A* formed the belief non-recklessly, the formation did not result from the exercise of a capacity to tell. Beliefs formed in such a way do not reflect *A*'s capacities as a knower, and therefore they do not deserve to be called "justifiably believed". Their sole merit is that *A* formed them while paying sufficient care to their veracity. Such care on *A*'s part indicates nothing about the belief's actual or probable veracity.

A case in which a judgment is merely rationally made was illustrated in section 3.3. To repeat: If I move away from the fens to live in bittern-ringer county, my capacity $C(\mathcal{T}\mathcal{H}\mathcal{A}\mathcal{T} - \mathcal{I}\mathcal{I} - A - \text{BITTERN}, \text{hearings of a bittern's booming})$ will no longer count as a capacity to tell. It will merely be a capacity to form beliefs. This does not mean that an exercise of this capacity must be reckless: Suppose I now hear the booming of one of the few actual bitterns in the county. Suppose further that I have no reason either to suspect that I am hearing a ringer or that there are such ringers about at all. In that case, if I exercise my capacity—thereby coming to believe, of the bittern, that it is a bittern—I will non-recklessly acquire a true belief. This belief will be rationally had. Still, since I will not acquire the belief through the exercise of a capacity to tell, I will not believe what I do justifiably. By living in bittern-ringer county, I robbed myself of my ability to tell bitterns by their booming. So the belief will be formed in a way that *guarantees* it does not amount to knowledge. And that is all the reason one needs to deny it the status of being justifiably believed.

3.5 A TEST CASE

SP has an number of attractive features which make it eminently plausible. First, it is highly intuitive as it is a mild rational reconstruction of our implicit conception of non-inferential perception based knowledge acquisition. Second, since it does not require that truth (veridicality, accuracy) conditions be assigned to experiences, it avoids the pitfalls of representational accounts of non-inferential experience based knowledge acquisition presented in chapter 2. Third, SP's relationalist conception of experiences leads to distinctively simple and compelling accounts of experiential tracking and experiential demonstrative concept use. Finally, SP's relationalist conception of experiences allows it to account for certain cases of non-inferential experience based knowledge, which embarrass a widely held view about the relationship between experiential phenomenology and experience based knowledge. This section will illustrate SP's last advantage.

Consider a certain line of thought concerning the relationship between experiential phenomenology and experience based knowledge. In arguing against what he calls the "Wallpaper View", Mark Johnston expressed it thus (2006, p. 263-264):

...suppose that when you look at the flag of Germany you enjoy qualia of just the same sort as I enjoy when I look at a similarly patterned flag whose constituent colors are green, blue, and white. Even though we have entirely different sensory deliverances, the Wallpaper View has the incredible consequence that we can both immediately know just by looking that a given German flag is red, yellow, and black. ...This can't be right. If external things are colored, then their colors must be tightly connected to the distinctive qualities which visual awareness alone reveals. Our immediate perceptual judgments about the colors of items in the environment do not systematically prescind from the qualities of surfaces ostensibly presented in sensory experience. Instead those judgments seem to predicate those very sensed qualities of environmental items. So it is hard to make clear how two subjects with inconsistent qualitative visual presentations could both be right in their immediate beliefs about the color of the relevant stripe in the flag.

According to Johnston, if some object looks different colors to A at t and at t' , then there cannot be a single color that A comes to (directly and on experiential grounds) know the object to be at both t and t' . The reason for this is that if something looks a certain color to A , A cannot come to (directly and on experiential grounds) know the thing to be any other color.

To show that Johnston’s line of thought must be rejected (and that SP can do better) I will revisit the Alice scenario from section 2.2. The scenario ran as follows:

At stage 1 Alice—someone with normal color vision—is 18 years old and broke. In need of funds, she agrees to undergo a red-green “Experiential Color spectrum Inversion” operation (an “ECI-op”, for short). During this operation, a chip is inserted into her optic nerve, where it replaces signals that pre-op were typically caused by red and reddish (green and greenish) things with those that pre-op were typically caused by green and greenish (red and reddish) things. The effect of the operation is this: Post-op, red and reddish things (in any lighting conditions C) look to Alice the colors that green and greenish things (in C) looked to her pre-op. Similarly, green and greenish things (in C) look to her the colors that red and reddish things (in C) looked to her pre-op.¹⁶ Before surgery, the doctors teach Alice about the ECI-op’s expected effect in great detail. In fact, they teach her, for any shade S (with its characteristic hue, saturation and lightness) that an object can have, and for any lighting conditions C , what shade things of shade S (in C) will look to her post-op.¹⁷

At stage 2 Alice wakes up from her ECI-op. Things look strangely colored to her, and she reports things like “the fire truck outside looks to me now the color fresh grass used to look to me”. Still, since the doctors taught her, for any S and C , how things of shade S (in C) will look to her post-op, she can tell what colors things are by sight. To illustrate how she performs this feat, consider a case in which she is faced with a red couch under good light. She reasons as follows: “This couch looks to me in this light the color green things looked to me in this light pre-op; and I know that in this light red things look to me the color green things used to look to me in this light pre-op. So the couch is red.” Similar forms of reasoning allow her to accurately tell by sight not only the hue of things, but also their color saturation, their lightness, and their level of warmth or cool. She can even tell which objects are similar and dissimilar in color. For example she can reason as follows: “Objects X, Y and Z look to me in this light the color green, red and cyan things (respectively) used to look to me in this light pre-op. I know that in this light red, green and purple things look to me the color green, red and cyan things (respectively) used to look to me in this light pre-op. So X, Y and Z must be red, green and purple (respectively). So Z is more similar in color to X than to Y.” Succinctly put: Post-op, Alice’s capacity to make color judgments—i.e., make judgments about the hues, saturation levels, lightness levels, warmth levels and color similarities of things—will be nearly indistinguishable from your pre-op capacity to make color judgments. Assuming, of course, that her reasoning skill and her acquaintance with the color space are good enough.

Why do I say *nearly* indistinguishable? Because even though her post-op sight-based color judgments will be so good that in the normal course of things they will never lead anyone to suspect she underwent an ECI-op, they will not be *identical* to what they were pre-op. The reason is that in normal human color vision there are more just noticeable differences

¹⁶I am presupposing that objects can truly be said to have colors (color properties). The presupposition, however, is not essential for my argument. Readers who cannot stomach it should feel free to replace my talk about things being colored with talk about things being spatially located, and my talk of a red-green spectrum inversions with talk of an upside-down location inversions.

¹⁷Imagine the doctors accomplish this using a large collection of Munsell color chips, matching chips that represent colored objects with chips that represent how those colors would look to her post-op in various lighting conditions.

(JNDs) of shade between red and blue than there are between blue and green (Kay, 1999; Palmer, 1999). There will therefore be distinct shades $S1$ and $S2$ such that pre-op Alice could just distinguish visually, but post-op she cannot; and there will be distinct shades $S3$ and $S4$ such that pre-op she could not distinguish visually, but post-op she just can. This fact might make itself manifest if she is asked to make sight-based color similarity judgments about determinate shades that are very close to each other - the sort of judgment that one typically does not make, but might be asked to make in a psychology lab.¹⁸

In the years after the ECI-op, Alice undertakes to make a few hundred sight-based color judgments a day. Over time, she finds that she can make subtler judgments faster, with less effort and sometimes without any conscious reasoning. By the time she is 50 (this is stage 3) this training regimen has long been abandoned. At stage 3, making sight-based color judgments is second nature to her. She makes them immediately, without any thinking or reasoning, and with effortless dexterity. In fact, the way she made sight-based color judgments before the ECI-op is identical, in all respects available to her consciousness, to the way she makes them now. Still, at stage 3 Alice remembers full well how things looked to her pre-op. She knows that fire trucks look to her the color fresh grass used to look to her. And as Taylor (1966) remarked, if she is asked to paint two pictures, one of how things looked to her pre-op and one of how they look to her post-op, she will do the following: In the post-op picture she will paint the grass with green paint. In the pre-op picture she will paint it with red paint.

At stage 4 Alice is 60 years old, and she has developed amnesia for the period up to age 50. In particular, she does not remember the following: (i) how things used to look to her before the ECI-op, (ii) that she so much a had the ECI-op, (iii) what the doctors taught her about how shades would look to her post-op, and (iv) that she ever arduously had to think before making color judgments. Alice also has no memories of episodes of remembering anything from the period up to age 50. Despite her amnesia, however, her color vision at stage 4 remains as it was at stage 2. There are still distinct shades $S1$ and $S2$ ($S3$ and $S4$) such that pre-op she could (not) just distinguish, but post-op she cannot (can). Furthermore, sight-based color judgments are still second nature to her, and she makes them in the same way and with the same accuracy as in stage 3.

I argued in section 2.2 that the Alice scenario is both conceivable and metaphysically possible. I also argued that if Alice were to look at a single well illuminated red couch at both stages 1 and 4, on both occasions Alice would (directly and experientially) judge the couch to be a single color - red. Finally, I argued that the epistemic statuses of Alice's judgments would be incompatible with certain versions of ER [e.g., McDowell's (1996)]. I will now further argue that, contrary to Johnston's view, both of Alice's (directly and experientially made) judgments would be knowledgeable. [Note: The arguments I am about to give closely resemble those made in section 2.2 . The weary reader might choose to skim them.]

¹⁸For a discussion of whether a behaviorally undetectable spectrum inversion scenario is empirically possible in humans, see Block (2007c); Broackes (2007); Hardin (1997); Harrison (1973); Hilbert and Kalderon (2000); Nida-Rümelin (1996); Palmer (1999).

I will provide three arguments that both of Alice's judgments are knowledgeable:

(1) I take it to be deeply intuitive that Alice's stage-1-judgment about the couch can be knowledgeable. Why? Well, largely for selfish reasons: I take it to be obvious that my own experience-based color judgments, as well as my judgments about the colors things look to me, can be knowledgeable. But there is no relevant difference between pre-op Alice and myself. In fact, the entire thought experiment could be run with me as its hero, instead of Alice. So it seems very hard for me to deny Alice the same status. And since the reader can go through the same selfish line of reasoning too, the reader should agree. Now: If Alice's stage-1-judgment is knowledgeable, how could a similar judgment fail to be knowledgeable if made at stage 2? After all, at stage 2 such a judgment would be *logically inferred* from (a) a judgment Alice makes about the color the couch looks to her - which would be knowledgeable because she would make it in just the way that she made similar judgments at stage 1, and (b) her knowledge (which she acquired through impeccable training), for any shade S and lighting conditions C , of what shade things of shade S (in C) look to her post-op. So at stage 2, Alice can come know, of the couch, that it is red as well. But, if she can come to know this at stage 2, how can she not be able to know this at stages 3 and 4? What sets stage 3 apart from stage 2 is that by stage 3 Alice has become an expert judge of color. The acquisition of such expertise, however, cannot rob Alice from having the ability to know what the expertise is an expertise to know. Similarly, what sets stage 4 apart from stage 3 is that Alice suffers from amnesia at stage 4. But the amnesia has no impact on Alice's expertise at making color judgments, just as it has no impact on any of Alice's other areas of expertise - making experience-based judgments about the locations of things, or about the identities of familiar people, for example. Since it has no impact on the expertise, it has no impact on Alice's ability to know, of the couch, that it is red (through the exercise of this expertise).

(2) There are a number of theoretical considerations supporting the view that Alice's judgments about the color of the couch (both at stage 1 and at stage 4) are knowledgeable.

(i) At stage 1, we may assume, all of Alice's color judgments are highly reliable. And it is part of the description of the case that at stages 2-4 Alice's capacity to make color judgments is nearly indistinguishable from her stage 1 capacity to make them. So at all stages, Alice's

color judgments are highly reliable. A reliabilist will therefore be forced to admit that Alice's judgments about the couch's color issue from a reliable belief-forming process. (ii) It follows from the description of the case that at stage 4 the way Alice makes sight-based color judgments is identical, in all respects available to her consciousness, to the way she made them at stage 1. So at both stages Alice makes judgments about the color of the couch that are experience-based, free of any inclinations to believe otherwise, free from contradictory beliefs, free from misgivings, and free of any other defeaters. So the internalist epistemologist should have no reason to object to the claim that both judgments are knowledgeable. (ii) Both at stages 1 and 4 Alice could offer the standard sorts of justifications we could give in support of our claim to experientially know the color of the couch. In particular, if asked why she believes, of the couch, that it is red, Alice could honestly and truthfully reply by uttering (on either occasion) "because I can tell a red thing when I see it", "because I see that it is red", or "because I see its being red".

(3) There is some empirical evidence that certain kinds of spectrum shifts and inversions are real, and even common. (I present them in section 2.2.) If these cases are real, then we are in the following situation: Even though actual people often agree in their experience-based color judgments, the same couch sometimes looks very different shades to them. But even when we know this fact, we are not moved in the slightest to doubt the knowledgeability of each other's color judgments. So, in fact, we do not take the color that something looks to a person to be a deciding factor in estimating that person's claim to know its color. And it seems highly implausible and deeply unmotivated to suppose, despite our actual inclinations, that the color that something looks to a person *is* a deciding factor in estimating this.

These arguments (in conjunction with those from section 2.2 showing that Alice's judgments are direct and have the appropriate content) should persuade the reader that Johnston's view should be rejected. But it remains to show that SP can do better than Johnston's view. In other words, it remains to show that SP can explain why Alice's judgments are knowledgeable. It is to this task that I now turn.

SP takes seriously the intuitive idea that if, on a particular occasion, *A* can (i) notice changes that *x* might undergo by using her sense organs, and (ii) attend to *x*, then *A* experiences *x* on that occasion. Since the performance of these feats (on an occasion) does

not require x to feel a particular way to A , it follows that A might experience x at both t and t' , even though x feels very different ways to A on these occasions. Thus, in particular, suppose red things (in any lighting condition C) do look to A at t the color that green things (in C) look to A at t' (and vice versa). It can still be the case that A looks at a red object at both t and t' , and that on both occasions A can notice changes the red color might undergo by using her eyes. And A might further be able to attend to the red color on both occasions. If so, then at both t and t' , A sees the object's color. It is just that the seen color looks different ways to A on these occasions.

This outlook allows SP to easily account for the possibility that at both t and t' A comes to know, of the object, that it is red. SP need only say that on both occasions A sees the object's being red (though these states of affairs look different ways to her), and that on both occasions A successfully exercises some capacity to tell $C(\mathcal{T}\mathcal{H}\mathcal{A}\mathcal{T} - \mathcal{I}\mathcal{S} - RED, \textit{seeings of something's being red})$. Both exercises would culminate in knowledge, of the object, that it is red. And in both cases, the knowledge would be had owing to a seeing of that object's being red. Finally, SP can say all this while allowing that the capacity A exercises at t is distinct from the capacity A exercises at t' .

3.6 CONCLUSION

The core of SP has now been presented and applied. Though the presentation of SP is not yet complete, the reader can already see that the view has the following advantages:

1. SP is a mild rational reconstruction of our implicit conception of non-inferential perception based knowledge acquisition.
2. SP accounts not just for perceptual knowledge, but also for perceptually based justified beliefs and for perceptually based rational beliefs.
3. SP accounts for learning on the strength of both on-line perceptual experiences, and remembered perceptual experiences.
4. SP does not require that truth (veridicality, accuracy) conditions be assigned to experiences, so it avoids the pitfalls of representational accounts of non-inferential experience

based knowledge acquisition presented in chapter 2.

5. SP offers distinctively simple and compelling accounts of perceptual tracking and perceptual demonstrative concept use.
6. SP can accommodate the fact that it is possible for people whose color spectra are inverted (relative to each other) to both knowledgeably and directly judge the colors of things.

In the remainder of this essay, I will argue that SP can provide an epistemology not just for perception based judgments, but also for judgments based on non-perceptual experiences. That SP can do so easily and in a unified way should count as a further advantage of the view.

Before turning to the epistemological powers of non-perceptual experiences, some preliminaries are in order. These will be addressed in chapter 4, which concerns the metaphysics of hallucinations. Chapter 5 will then return to the topic of SP.

4.0 OBJECTS OF HALLUCINATION

In a critical scene of Shakespeare's *Macbeth*, the eponymous character hallucinates a dagger and wonders:

Is this a dagger which I see before me,
The handle toward my hand? Come, let me clutch thee:
I have thee not, and yet I see thee still.
Art thou not, fatal vision, sensible
To feeling as to sight? Or art thou but
A dagger of the mind, a false creation,
Proceeding from the heat-oppresed brain?

At first blush, Macbeth's predicament seems perfectly easy to explain: He is hallucinating a certain dagger, and wondering whether it is a real dagger (i.e., the sort of thing that might stab someone), or an unreal dagger (i.e., a mere figment of his mind). He tries to test which of these hypotheses is correct by trying to clutch this dagger. His failure gives him some indication that the dagger is unreal.

The explanation just given, however, presupposes that there is something that Macbeth is trying to clutch and about whose reality he is wondering. So if we accept this explanation, we will be saying that there are such things as hallucinated daggers. Should we really say this? I will use the name "realism" for the view that we should; that there indeed are hallucinated entities. But realism typically does not just *prescribe* that we ontologically commit ourselves to hallucinated entities. It also claims to *describe* our preexisting ontological commitment to such entities. In other words, realism typically incorporates the view that when we think or assert things like "there is a hallucinated dagger before me", we demonstrate that we already hold that hallucinated entities exist.

At first blush, however, realism (with or without its descriptive part) seems suspect. Isn't it just obvious that we should not be in the business of saying that there are hallucinated entities? Aren't hallucinated entities just as clearly non-existent as the present king of France or square circles? And even if we do not find hallucinated entities essentially odious, don't their costs outweigh their benefits? Aren't they incompatible with a physicalist worldview?

I will use the name "irrealism" for the view that we should deny that there are hallucinated entities. According to irrealists, the verb "to hallucinate" has an intensional transitive use, on which it can be true that Macbeth hallucinates a dagger, even though what Macbeth hallucinates does not exist. Irrealists further say that in asking "art thou but a dagger of the mind, a false creation?", Macbeth is entertaining a referentially incomplete thought, or a mock-thought, in which the mental indexical "thou" fails to refer.

All forms of irrealism have an immediate question to answer. To see it, consider a few potential answers Macbeth might give to his "art thou but a dagger of the mind"-question:

- (1a) Thou art a mere dagger of the mind.
- (1b) Thou art not a real dagger.
- (2a) Thou art not a dagger of the mind.
- (2b) Thou art a real dagger.

There is a strong feeling that were Macbeth to answer his own question with either (2a) or (2b), he would *not* give the question a correct answer. Judgments (2a) and (2b) strike us as false (or, as neither true nor false). Equally, there is a strong feeling that were Macbeth to answer the question with either judgment (1a) or (1b), he *would* give it a correct answer. Judgments (1a) and (1b) strike us as true. And, what is perhaps clearer, there is a strong feeling that judgments (1a) and (1b) have a different truth value than judgments (2a) and (2b). But, if we are to say that the mental indexicals "thou" occurring in (1a) and (1b) fail to refer, we must also say that (1a) and (1b) are as untrue as (2a) and (2b). This seems wrong, and contrary to strong intuitions. We feel we can both "get things right" and "get things wrong" in accepting thoughts that appear to commit us to the existence of hallucinated entities. As irrealists, are we to simply ignore these feelings?

Revisionary irrealism answers this question in the positive. According to revisionary irrealists, many of us bear the attitude of belief to thoughts which, like (1a) and (1b), can only be true if there are hallucinated entities. Thus, revisionary irrealists accept the

descriptive part of the typical realist view. Revisionary irrealists only add to this, that we should disabuse ourselves of this ontological commitment. We should thus accept that (1a) and (1b) are as untrue as (2a) and (2b), despite our intuitions. If we'd like, we might replace believing things that require the existence of hallucinated entities with believing things less ontologically committed. But the crucial point is that this will require a significant change in what we believe.

Against revisionary irrealism stands hermeneutic irrealism. Hermeneutic irrealists deny that we (or, at least, those of us who are not card-carrying realists) need to revise our intuitions regarding (1a)-(2b). This is not because they think (1a) or (1b) are true. Like all irrealists, hermeneutic irrealists encourage us to deny this. Rather, it is because they do not think we strictly speaking believe (1a) or (1b). Strictly speaking, we do not bear the attitude of belief to *any* thoughts whose truth requires the existence of hallucinated entities. Nor do we find it intuitive that such thoughts are literally true. So by hermeneutic irrealist lights, all but card-carrying realists handle such thoughts in a different way. And different versions of hermeneutic irrealism specify what that way is differently.

This chapter constitutes a sustained argument for a particular version of realism, and against both revisionary and hermeneutic irrealisms. The argument proceeds as follows: After a short preamble (section 4.1.1), section 4.1.2 reviews the various versions of hermeneutic irrealism. Sections 4.1.3-4.1.5 then put forward three considerations against these versions of the view. A possible hermeneutic reply is explored and rejected in section 4.1.6. With hermeneutic irrealism thus out of the way, section 4.2.1 presents the recommended version of realism - abstract realism. It claims that abstract realism best captures folk intuitions regarding hallucinated entities. Section 4.2.2 then discusses the nonexistence objection to this claim, according to which the folk intuitively hold that hallucinated entities do not exist. In rejecting this objection, the section criticizes intentional inexistence views (also known as Meinongian views). This leaves abstract realism only revisionary irrealism to contend with. Sections 4.3.1-4.3.4 therefore showcase intuitive, phenomenal, epistemological and semantic advantages that abstract realism has over revisionary irrealism. Though these sections do not deliver a knock-down objection to the revisionary position, they do show that it falls far behind the theoretical elegance of its realist rival. The chapter closes with section 4.4, which

answers a number of common objections to realism.

Though primarily an argument for abstract realism, this chapter is also a plea for an interdisciplinary look at hallucinations. On the one hand, hallucinated objects are of interest to metaphysicians studying other unusual entities, such as fictional objects or numbers. On the other hand, hallucinations are of interest to philosophers of mind studying the nature, phenomenology and epistemology of other experiential states, such as perceptions and illusions. In this way, hallucinations bring together philosophical problems that are more commonly studied in isolation. Sadly, however, current philosophical literature mostly ignores this fact. Metaphysicians working on hallucinated objects frequently do not explore what philosophers of mind say about the nature of hallucinations. Similarly, philosophers of mind frequently do not take seriously the possibility that hallucinated objects exist.¹

One way to read this chapter is as an attempt to bring both strands of the literature together. While sections 4.1 and 4.2 are mainly devoted to the traditional metaphysical discussion, sections 4.3 and 4.4 bring philosophy of mind considerations to the forefront. I suspect that this makes the case for hallucinated objects stronger than the case for other unusual objects.

4.1 TROUBLES FOR HERMENEUTIC IRREALISM

4.1.1 A First Charge for Realism

Suppose Macbeth (who, I remind the reader, is *not* a piece of mere fiction, but a real man who actually was the king of the Scots between 1040 and 1057) is hallucinating a scene in which some witch, Morgana, creates some sharp dagger, Dag. Suppose further that he responds to this hallucination by asserting:

(3) Morgana created Dag.

There seems to be something right about (3). But is (3) true? The question seems puzzling because there are two ways in which (3) might be evaluated.

¹Johnston (2004) and Smith (2002) are notable exceptions to this.

First, (3) might be valuated for truth simpliciter, i.e., it might be evaluated as a simple report true iff Morgana created Dag. It would be appropriate to evaluate (3) in this way if Macbeth asserted (3) without realizing he is hallucinating, and upon being taken in by his hallucination. When evaluated for truth, (3) turns out to be untrue on both realist and irrealist lights. Irrealists would say this report is untrue because “Morgana” and “Dag” do not refer. Abstract realists would say that (3) is false, because though Dag and Morgana do exist, it is not the case that Morgana created Dag. Dag - like Morgana herself - was actually created by Macbeth, who hallucinated the two into existence.

But there is a second way of evaluating (3). We might evaluate (3) not for truth simpliciter, but for truth in the relevant hallucination. It would be appropriate to evaluate (3) in this way if Macbeth asserted (3) upon realizing that he is hallucinating and in the context of an attempt to report on what he is hallucinating. And there is wide agreement between realists and irrealists that (3) is true in Macbeth’s hallucination.

The distinction between evaluating an assertion for truth and evaluating it for truth in the relevant hallucination makes life a bit hard for the typical realist. The typical realist would like to argue that the folk are committed to the existence of entities such as Dag and Morgana, because they often take assertions like (3) to be true. But the hermeneutic irrealist’s response is to say that the folk don’t take assertions such as (3) to be true - they merely take them to be true in the relevant hallucination. So the folk are not committed to the existence of creatures such as Morgana and Dag. They are merely committed to holding that they exist in the relevant hallucination.

Realists have a standard response, however. Following the work of Saul Kripke (2013), Nathan Salmon (1998), Stephen Schiffer (1996), John Searle (1975), Amie Thomasson (1999), Peter Van Inwagen (1977) and others, realists point the hermeneutic irrealist’s attention to assertions such that (i) the folk clearly take them to be true in some sense, (ii) it is highly implausible that the folk take them to be true in the relevant hallucinations, and (iii) the truth of the assertions requires the existence of hallucinated entities. Let us call such assertions “critical assertions”. The realist’s argument is that the folk do take critical assertions as true (simpliciter), and that they are therefore committed to the existence of hallucinated entities.

As our example of a critical assertion, let us pick:

(4) Macbeth created Dag.

There is a strong intuition that (4) is true. If Macbeth were duped by his hallucination and asserted (3), (4) is what we would say to set him straight. “You’re wrong”, we’d tell Macbeth, “Dag is something you created, a creature of your mind. It is not a real dagger made by a real witch.” At the same time, (4) is clearly *false* in Macbeth’s hallucination. In the hallucination, it is Morgana—and not Macbeth—that created Dag. So the realist’s charge is that our intuition that (4) is in some sense true cannot be explained as the intuition that (4) is true in Macbeth’s hallucination. Rather, the intuition that (4) is in some sense true is just the intuition that (4) is true simpliciter. And if we feel that (4) is true simpliciter, we are committed to Dag’s existence.²

We need not have used (4) as our example of a critical assertion. Here are a few other examples of critical assertions that would have done just as well: “Dag is a merely hallucinated dagger”, “Dag is a figment of Macbeth’s mind”, “Dag is not a real dagger”, “Dag is distinct from this dagger” [asserted while pointing to a real dagger], “Dag symbolizes murder to Macbeth”, “Lady Macbeth told everyone about Dag”, “Dag is famous”, “Macbeth worships Dag”, “Dag terrifies Macbeth” and “Dag makes Macbeth’s blood pressure rise”. If we fill in some more details about the circumstances surrounding (and the effects of) Macbeth’s hallucinations of Dag, we can make all these assertions sound true. At the same time, we can also make none of these assertions sound true in Macbeth’s hallucinations. And finally, the truth of these assertions requires Dag’s existence.

The pressure critical assertions put on hermeneutic irrealism has not gone unnoticed. Three versions of the view offer three distinct responses to it.

²I am using (4) as my example since “creation” sentences are widely and standardly cited in the context of the realism debate. But it has been suggested to me that (4) might sound odd to those supposing that Macbeth did not intentionally set out to hallucinate Dag. Readers who share this concern may replace “Macbeth” in (4) with “Macbeth’s mind” or “Macbeth’s visual system”. Further examples will be given in the next paragraph. But I would also like to note that I do not share the concern about (4). Just as “Macbeth created a noise” would sound true even if he did so by unintentionally falling down the stairs, “Macbeth created Dag” should sound true even if he did so by unintentionally undergoing a hallucination.

4.1.2 Three Hermeneutic Irrealisms

Hermeneutic irrealists deny that the folk (i.e., those of us who are not card-carrying-realists) believe the literal content of their critical assertions. When the folk make a critical assertion, they are not expressing a belief with the literal content of the assertion. They are doing something else instead. Three different kinds of hermeneutic irrealism—operator irrealism, paraphrase irrealism and attitude irrealism—correspond to three different ways of specifying what that other thing is.

Let “ p ” stand for any critical assertion. According to operator irrealists, when the folk assert “ p ”, they are expressing their belief that $O(p)$, where “ O ” is some prefixing operator. For example, operator irrealists would say that when Macbeth says “I created Dag”, what he actually believes is that *according to the realist hypothesis*, he created Dag. Regardless of what the prefix operator “ O ” is, operator realists insist that believing that $O(p)$ does not commit one to the existence of hallucinated entities, even if believing that p does. Furthermore, operator irrealists give the following account of our intuitions about critical assertions: Intuitions which seem to be intuitions that these assertions are true are in fact intuitions that these assertions, when prefixed with the relevant operator, are true.

Paraphrase irrealism is a close cousin of operator irrealism. Like operator irrealists, paraphrase irrealists hold that when the folk assert “ p ”, they are expressing some belief other than that p . But where operator irrealists specify what that other belief is by manipulating “ p ”, paraphrase irrealists take a different approach. They think that when the folk assert “ p ”, what they actually believe is that q , where “ q ” (i) need not be a modification of “ p ”, (ii) need not involve a prefixing operator, and (iii) can be true even if there are no hallucinated entities. In this sense, “ q ” is an irrealist-friendly paraphrase of “ p ”. For example, paraphrase irrealists could say that when it seems that Macbeth believes that he created Dag, what he actually believes is that he created the hallucination according to which Dag exists. Like their operator irrealist counterparts, paraphrase irrealists say that intuitions which seem to be intuitions that critical assertions are true are in fact intuitions that their irrealist-friendly paraphrases are true.

Unfortunately, neither operator- nor paraphrase-irrealism has recently been championed

in debates on hallucinations. However, there has been significant related work done with regard to discourse whose truth requires the existence of *fictional* entities (e.g. Sherlock Holmes, Frankenstein’s monster). In particular, Stuart Brock (2002, p. 9) has suggested that “Mary Shelley created Frankenstein’s monster” can be replaced by “according to the realist hypothesis, Mary Shelley created Frankenstein’s monster”. Mark Sainsbury (2010, p. 107) has suggested to replace the same with “Mary Shelley created the story according to which Frankenstein’s monster is a monster”. It is not clear (at least to me) whether Brock or Sainsbury hold that their replacements capture the content of beliefs the folk express in asserting “Mary Shelley created Frankenstein’s monster”. I therefore mean to suggest *neither* that Brock holds a view akin to operator irrealism for the realm of fictional entities, *nor* that Sainsbury is friendly to a counterpart of paraphrase realism for fictional entities. However, since their suggested replacements can easily be transformed into replacements for assertions seemingly about hallucinated entities, I will draw on their insightful work in discussing both operator- and paraphrase irrealism.

The alternative to both operator and paraphrase irrealism is attitude irrealism. According to this view, when the folk assert “*p*”, their assertion does not express a belief at all. Rather, it expresses the *make-belief* (alternatively: the pretense, simulation, imagining) that *p*. While believing that *p* commits the believer to the existence of the entities required for its being the case that *p*, make-believing that *p* involves no such commitment.

Attitude irrealists further hold that when we make-believe that *p*, we do so in an attempt to be faithful to a relevant fiction. What is the case according to this fiction is determined by certain rules (called “principles of generation”). Attempting to be faithful to this fiction is thus attempting to make-believe just what the rules determine is the case according to the fiction. For example, given Macbeth’s hallucinatory circumstances, the rules for a certain realist-inspired fiction determine that, according to this fiction, Macbeth created Dag. Assertions of (4), which are made in the course of an attempt to be faithful to this fiction, are therefore true according to it. And the decisive point for the attitude irrealist is that the existence of the fiction makes it appropriate (perhaps even required) to evaluate critical statements that attempt to be faithful to it not for truth simpliciter, but for truth according to that fiction. It is in this way that attitude irrealism attempts to vindicate our intuitions

regarding critical assertions: Intuitions which seem to be intuitions that critical assertions are true are in fact intuitions that they are true according to the relevant fiction.

The clearest champion of attitude irrealism is Gareth Evans (1982, chapter 10).³ Gregory Currie (1990), Anthony Everett (2005; 2013), Frederick Kroon (2011), Kendall Walton (1990) and others hold views akin to attitude irrealism for the realm of fictional entities. I will draw on the insightful work of all of these thinkers in discussing attitude irrealism.

I mean to make trouble for hermeneutic irrealism. I will raise three problems that apply generally to all three versions of the view.

4.1.3 The Problem of Ignorance and Error

According to hermeneutic irrealism, when the folk assert “ p ” (where “ p ” is a critical assertion), they are doing something other than express their belief that p . But hermeneutic irrealist specifications of what it is that the folk are doing instead all require attributing to them various beliefs (or concepts), which they may lack due to either ignorance or error. So, under conditions of ignorance or error, the folk might make critical assertions perfectly reasonably, though a hermeneutic irrealist analysis of their assertion cannot be given. And this means that hermeneutic irrealism is untenable, at least in ignorance and error cases.

Consider (4) again. An ignorance case with respect to (4) is given by Irma, who honestly and reasonably asserts: “I have heard so many conflicting reports, that I take no stand on whether Macbeth ever hallucinated anything. I also take no firm stand on what Dag is exactly. Dag is either a real dagger, a merely hallucinated one, a merely fictional one or a dagger drawing. Still, there is one thing all the evidence points to and which I do believe: Macbeth created Dag.” An error case with respect to (4) is given by Ed, who honestly and reasonably asserts: “I have all the following on excellent authority: Dag is a merely fictional dagger. Macbeth created Dag. Macbeth created Dag by writing a fictional novel in which Dag appears. And, by the way, Macbeth never had a hallucination.” It will now be shown that hermeneutic irrealist analyses of what Irma and Ed are doing in asserting (4) cannot be given.

³Note, however, that Evans mostly focused on assertions such as (3), rather than on critical assertions.

Start with operator irrealism. As we saw, an operator irrealist inspired by Brock (2002) would suggest that when the folk assert (4), they in fact express the belief that

(4a) According to the realist hypothesis, Macbeth created Dag.

Alternatively, an operator irrealist might suggest they express the belief that

(4b) According to the contextually relevant hypothesis, Macbeth created Dag.

Or perhaps they express the belief that

(4c) According to the hypothesis that r , Macbeth created Dag,

where “ r ” is some specification of content which logically entails all and only those critical assertions that are intuitively true. (For example, a part of “ r ” might be “there are merely hallucinated entities, and Dag is one of them, and merely hallucinated entities can be created, and Macbeth created Dag, and ...”.)

None of these can be what Irma and Ed believe. Irma and Ed do not believe that Dag has anything to do with Macbeth’s (or anyone else’s) hallucinations. So they do not believe that it is according to any hypothesis concerning the existence and properties of hallucinated entities that Macbeth created Dag. So they do not believe (4a)-(4c). Besides, Irma and Ed might not believe (4a)-(4c) for plenty of other reasons. Believing (4a) requires the believer to refer to the realist hypothesis using the concept “the realist hypothesis”. But Irma might not have grasped this concept, and Ed might mistakenly believe that the hypothesis according to which Macbeth created Dag is rather the verificationist one. These shortcomings are compatible with Irma and Ed’s honestly and reasonably asserting (4). Similarly, believing (4b) requires the believer to have views about what is the case according to the contextually relevant hypothesis. But Irma might have no idea how to determine what is the case according to the contextually relevant hypothesis. And Ed might think that according to this hypothesis, Dag has no creator. Again, these shortcomings are compatible with Irma and Ed’s honestly and reasonably asserting (4). Finally, believing (4c) requires making explicit the content of a hypothesis that logically entails all and only those critical assertions that are intuitively true. So doing requires, apart from significant insight, exhaustive knowledge of the circumstances of all hallucinations. Such insight and knowledge are, to put it mildly, not very common

among the folk. And nonetheless the folk can go on honestly and reasonably making critical assertions.⁴

Consider paraphrase irrealism next. Paraphrase irrealists drawing on Sainsbury (2010, p. 107) would suggest that by asserting (4) the folk express the belief that

(4d) Macbeth created the hallucination according to which Dag exists.

But Irma and Ed do not believe (4d). They do not believe it, because they do not even believe that Macbeth ever had a hallucination. But that does not make their assertions of (4) any less honest or reasonable.⁵

This brings me to attitude irrealism. By the lights of this view, when the folk assert (4) they express the *make*-belief that Macbeth created Dag. They undertake this make-belief in the course of an attempt to make believe what certain rules determine to be the content of some fiction, *f*. An attitude irrealist familiar with Walton's (1990, p. 416-424) and more particularly with Everett's (2013, p. 59) would suggest that these rules include:

(R) According to *f*, *x* created *y* iff there is some hallucination *h* such that: (i) according to *h*, *y* exists; (ii) *x* had hallucination *h*; (iii) there is no hallucination *i* earlier than *h* such that according to *i*, *y* exists; and (iv) “*y* exists” is not true.

So by attitude irrealist lights, folks who assert (4) are attempting to make believe in accordance with (R). But if so, asserters of (4) had better believe (or presuppose, or otherwise take it) that all the following are the case: (a) Macbeth had a certain hallucination, *h*, (b) according to *h*, Dag exists, (c) there is no hallucination *i* earlier than *h* such that according to *i*, Dag exists, and (d) “Dag exists” is not true. And the problem with this suggestion is

⁴Two extra worries regarding (4c): First, since “*r*” must be quite intricate, expressing a belief such as (4c) in asserting (4) will tax the asserter's memory and attentional resources. But we witness no such attentional and memory burdens among those who assert (4). Second, a belief such as (4c) will be *a priori*; whereas an assertion of (4) intuitively expresses an *a posteriori* truth. The reason (4c) will be *a priori* is that “*r*” logically entails that Macbeth created Dag. The belief the folk express in asserting (4), however, is a belief the folk cannot gain from the armchair.

⁵I should add that (4d) is not a proper paraphrase for (4) for another reason: Macbeth can hallucinate ordinary things as well as merely hallucinated ones. So suppose Macbeth hallucinates not Dag but a real person - Lady Macbeth, say. In this case, though it is true that Macbeth created the hallucination according to which Lady Macbeth exists, it is false that Macbeth created Lady Macbeth. (The lady, after all, was created by her parents.) This shows that (4) should be paraphrased some other way, perhaps as (4e): “Macbeth created the hallucination according to which Dag exists, and ‘Dag’ does not refer.” Clearly (4e) is no less vulnerable than (4d) to the objection posed by the Irma and Ed cases.

that Irma and Ed neither believe nor presuppose (a), (b) or (c), and yet they are perfectly reasonable and honest in asserting (4).

4.1.4 The Problem of Communication

Realism, even if false, is surely a coherent view. So even if there are in fact no realists, there could have been. There could even have been English speaking realists. And such realists would express their beliefs about hallucinated entities by making critical assertions in English. In their mouths, “Macbeth created Dag” would express the belief that Macbeth created Dag, and similarly for the other critical assertions.

Now suppose Alice is such a realist, and suppose Bob is just one of the folk. The two meet, and have the following conversation:

Bob: “Macbeth had a hallucination in 1040. He hallucinated some dagger, Dag, right in front of him. He never hallucinated Dag before, and Dag isn’t one of the daggers he ever perceived.”

Alice: “Aha. So in 1040 Macbeth created the merely hallucinated dagger Dag.”

Bob: “Yes, that’s precisely what I think.”

This seems to be a case of perfectly successful communication, culminating in both Alice and Bob’s believing what Alice expresses by saying “So in 1040 Macbeth created the merely hallucinated dagger Dag.” In fact, it seems that both Alice and Bob believe this on the same grounds - those that Bob gave in his opening assertion.

The case creates a dilemma for the hermeneutic irrealist: Either the conversation culminates in Alice and Bob’s believing the same thing (which Alice expressed by saying “So in 1040 Macbeth created the merely hallucinated dagger Dag”), or not. If they do believe the same thing, hermeneutic irrealism is undermined. Since Alice is a realist, the belief she expressed entails that merely hallucinated entities exist. If Bob believes the same thing, his belief also entails this - contradicting hermeneutic irrealism. Alternatively, if Alice and Bob do not believe the same thing, then their communication is unsuccessful. But if so, why do all appearances suggest otherwise?

This dilemma may fail to impress some hermeneutic irrealists. Such irrealists may accept the second horn and hold that Alice and Bob are miscommunicating. After all, they will say,

Alice is committed to the existence of hallucinated entities, whereas Bob - like the rest of the folk - is not. It is for this reason that the two do not end their conversation by sharing the same belief. The fact that this is not immediately evident from their communication is neither here nor there.

In response to such intransigence, I would like to strengthen the argument from communication. The strengthened argument appeals to the ignorance and error cases familiar from the previous section (4.1.3). Its overall structure is as follows: (i) Successful communication cases can show that realists and “ignorant” folk express the same beliefs by their critical assertions. (ii) Successful communication cases can also show that well-informed folk and “ignorant” folk express the same beliefs by their critical assertions. (iii) Therefore, realists and well-informed folk express the same beliefs by their critical assertions. And since realists are committed (by their critical assertions) to the the existence of hallucinated entities, it follows from (iii) that (iv) well-informed folk are committed (by their critical assertions) to the the existence of hallucinated entities. And (iv) is incompatible with hermeneutic irrealism.

It remains to establish premises (i) and (ii). For (i), suppose Alice and Irma meet and have the following conversation:

Irma: “Macbeth created Dag. I know that well enough. But do you know when he created Dag?”

Alice: “Yes. Macbeth created Dag in 1040. He told me so himself.”

Irma: “Aha. So we are agreed: Macbeth created Dag in 1040.”

This again seems to be a case of perfectly successful communication, culminating in Irma and Alice’s believing the same thing on the strength of Macbeth’s testimony. And, since what Alice believes is that Macbeth created Dag in 1040, it seems that so does Irma. What is special about having Irma (rather than Bob) be part of the conversation, is that it is going to be very difficult for hermeneutic irrealists to maintain that Irma and Alice are miscommunicating. Irma, recall, is ignorant of Dag’s being a merely hallucinated entity and of Macbeth’s having hallucinations. By the considerations of the previous section, it follows from this that her assertion, “Macbeth created Dag in 1040”, does not express what hermeneutic irrealists suggest it does (i.e., some other belief, or a make-belief). So it must be that her assertion expresses the belief that Macbeth created Dag in 1040.

For (ii), suppose Bob meets with Irma. Suppose further that they have the following conversation:

Irma: “Macbeth created Dag. I know that well enough. But do you know when he created Dag?”

Bob: “Yes. Macbeth created Dag in 1040. He told me so himself.”

Irma: “Aha. So we are agreed: Macbeth created Dag in 1040.”

Here too, it seems that the conversation successfully culminates in Irma and Bob’s being in the same mental state (be it a single belief or a single make-belief) on the strength of Macbeth’s testimony. Could it be plausible for a hermeneutic irrealist to nonetheless insist that Irma and Bob actually end up in distinct mental states? No. The only reason it was plausible for hermeneutic irrealists to hold that Alice and Bob end up in distinct mental states was that Alice and Bob had distinct views regarding the existence of hallucinated entities. But Irma and Bob do not have such distinct views. Irma is just as much one of the folk as Bob. Neither is a card-carrying realist. So a hermeneutic irrealist should take Irma and Bob’s critical assertions to express the same mental state. An even stronger point is this: The only thing that sets Irma and Bob apart, is that Irma has a few fewer beliefs than Bob. Such ignorance cannot constitute grounds of attributing ontological commitments to Irma that are not also attributed to Bob. So it is implausible to hold that Irma’s assertion of “Macbeth created Dag in 1040” commits her to Dag’s existence, whereas Bob’s assertion of the same does not.

With premises (i) and (ii) thus supported, the rest of the argument is clear: By the argument for (i), Irma’s assertion of “Macbeth created Dag in 1040” commits her to Dag’s existence. But this entails, given (ii), that Bob’s assertion of the same also commits him to Dag’s existence. And this shows that hermeneutic irrealists analyses of Bob’s critical assertion are untenable.

4.1.5 The Problem of Theoretical and Practical Inference

If “ p ” is a critical assertion, hermeneutic irrealists deny that folk who are sincerely willing to assert “ p ” believe that p . But if that is correct, such folk (assuming they are rational) would not behave as only those who do believe that “ p ” should behave. And the problem before us

is that rational folk often do behave as only those who do believe that “*p*” should behave. The behavior I have in mind is the drawing of certain theoretical and practical inferences.

Let’s start with theoretical inferences. Suppose Bob from the previous section is trying to figure out if Macbeth created the *real* dagger, Ger, at time T. He answers his own question by reasoning (out loud) thus:

- (5) All Macbeth created at T is Dag.
 - (6) Dag is not real, but Ger is.
- Therefore,
- (7) Macbeth did not create Ger at T.

After reasoning in this way, Bob will in every way behave as one who believes (7). This suggests that the reasoning culminates in Bob’s believing (7). Furthermore, since (7)’s truth does not require the existence of hallucinated entities, (7) is not a critical assertion. So there is no hope for the suggestion that Bob’s reasoning culminates in his believing something other than (7) or in his merely make-believing (7).

Given that Bob’s inference culminates in his believing (7), we must ask how Bob forms this belief. The answer is clear: Bob forms this belief *rationally*, by inferring it from (5) and (6). And now trouble begins for the hermeneutic irrealist. If Bob rationally comes to believe (7) by inferring it from (5) and (6), Bob believes (5) and (6). But (5) and (6) are critical assertions and so, by hermeneutic irrealist lights, Bob does not believe (5) and (6). Therefore, hermeneutic irrealism is false. To see this in more detail, note that according to attitude irrealism, (5) and (6) express make-beliefs. But forming a belief such as (7) on the strength of make-beliefs is irrational; and Bob is not irrational. Alternatively, according to operator and paraphrase irrealisms, (5) and (6) express beliefs with contents distinct from the literal contents of (5) and (6). But it is easy to verify that the beliefs these versions of irrealism would have (5) and (6) express are no grounds for forming the belief (7) expresses.

Now let’s consider practical inferences. Suppose Bob intends and desires to write the complete catalog of Macbeth’s creations. He already embarked on this project, and his manuscript contains entries for each of Macbeth’s biological children, for the cities he built, for the documents he authored, and even for that hallucination he had in 1040. Still, the catalog does not have an entry for Dag. But this changes one day. On that day Bob asserts (4) and immediately proceeds, driven by his desire for historical exhaustiveness, to add a

catalog entry for Dag.

Bob's behavior is precisely the behavior we would expect from someone with a keen sense for historical accuracy who is also willing to assert things like (4). His behavior is therefore perfectly rational. But if Bob's behavior is rational, then Bob believes (4). Without this belief, Bob's desire to write the catalog could not lead - through practical inference - to the behavior of adding the Dag-entry. Besides, if Bob added the entry without believing (4), Bob would risk writing an inaccurate catalog. And Bob desperately wants to avoid that. So we can conclude that Bob believes (4), contrary to hermeneutic irrealism. To see this in more detail, note that according to attitude irrealism, Bob only make-believes that Macbeth created Dag. But to intentionally add the Dag-entry to the manuscript on the strength of a mere make-belief is irrational; and Bob is not irrational. Alternatively, according to operator and paraphrase irrealisms, Bob believes (4a), (4b), (4c) or (4d). But it is easy to verify that these beliefs cannot serve as grounds for intentionally adding the Dag-entry to the manuscript.

Attitude irrealists might try to resist the last criticism. They might suggest that Bob's behavior constitutes not an intentional adding of a Dag-entry to the manuscript, but an intentional *pretended* adding of a Dag-entry to the manuscript. And they might further suggest that this pretended behavior is the rational result of (i) Bob's *pretended* belief (i.e., his make-belief) that Macbeth created Dag, and (ii) Bob's *pretended* desire (i.e., his make-desire) to write the complete catalog of Macbeth's creations. But this cannot be right. First, Bob does not pretend to desire to write the complete catalog of Macbeth's creations. He actually desires to do so. If Bob merely pretended to desire to write the catalog, he would not end up producing a detailed manuscript of it. Second, Bob's behavior does not constitute a pretended adding of a Dag-entry to the manuscript. Bob really does add a Dag-entry to the manuscript. At no point will Bob erase the entry, or give any other indication that his action was not to be taken seriously. So to regard his behavior as part of a pretense is *ad hoc* and implausible.

4.1.6 A Reply?

I think the above three objections are excellent reasons to hold that hermeneutic irrealism is false. But before setting this view aside, I would like to briefly consider a reply that might lead one to hope the view can be saved.

The reply can be put as follows: All that the above three objections show is that when the folk make the critical assertion “ p ”, they express their belief that p . But the objections do not show that the folk thereby undertake a commitment to the existence of hallucinated entities. Rather, in making the objections it was simply assumed that folk who fully believe that p (where “ p ” is a critical assertion) are automatically committed to the existence of hallucinated entities. This assumption can be resisted. It can be maintained even in fully believing what a critical assertion expresses, the folk believe something whose truth does not require the existence of hallucinated entities. And if the truth conditions of the belief that p are such that it can be the case that p even though hallucinated entities do not exist, then the folk may both believe that p and not be committed to the existence of hallucinated entities.

Let us examine this reply. Its plausibility seems to rest on the plausibility of assigning ontologically “non-committal” truth conditions to the beliefs that critical assertions express. So let us see how plausible it is, e.g., to assign the belief that Macbeth created Dag truth conditions that allow it to be true even when Dag fails to exist. I will consider two broad ways of making such an assignment.

The first kind of irrealist proposal might be to assign the belief that Macbeth created Dag the truth conditions of (4a), (4b), (4c), (4d) or the appropriate instance of (R). On this proposal, the folk may happily believe that Macbeth created Dag, since the belief commits them not to Dag’s existence, but merely to the existence of hallucinations, hypotheses or fictions with certain features.

The trouble with this proposal can be most readily seen if we compare the belief that Macbeth created Dag with the belief that Macbeth created Ger. Ger, recall, is an ordinary real dagger. So it should be clear that the truth of the latter belief does require Ger’s existence. Furthermore, it should be clear that the truth of the latter belief does not require

the existence of any hallucinations, hypotheses or fictions with certain features. Therefore, to hold that the *former* belief has the proposed truth conditions, one must also hold that there is a dis-analogy between the former belief and the latter one. Though both are beliefs in which the concept “ x created y ” is flanked by simple concepts purporting to refer to particulars, this similarity is not reflected in their truth conditions. And this leads to three problems.

First, it is widely accepted that our ability to entertain an endless number of beliefs (the “productivity” of belief), and the systematic relations between the beliefs we have and those we can entertain (the “systematicity” of belief), show that the truth conditions of our beliefs are compositionally determined by their logical form and by the concepts that occur in them (the “compositionality” of belief). Now apply this compositionality principle to beliefs in which the concept “ x created y ” is flanked by simple concepts purporting to refer to particulars. It emerges that all such beliefs are true iff the single relation denoted by the concept “ x created y ” obtains between the reference of the subject-position concept and the reference of the object-position concept. Unfortunately, the current irrealist proposal violates this. On this proposal, the belief that Macbeth created Dag requires for its truth the existence of hallucinations, hypotheses or fictions with certain features. The belief that Macbeth created Ger requires no such thing. So there can be no single relation denoted by the concept “ x created y ” on which the truth of *both* beliefs turns. As this violates the compositionality of belief, the current irrealist proposal must be rejected.

A second problem with the current proposal is that it entails that inferences that proceed in accordance with good formal deductive rules of reasoning may fail to be truth preserving. Consider, for example, the inference:

- (4) Macbeth created Dag.
- (6a) Dag is identical to Ger.
- Therefore,
- (4e) Macbeth created Ger.

A person who comes to believe that (4e) by inferring it from the beliefs that (4) and that (6a) is reasoning in accordance with the Indiscernibility of Identicals. But despite the rule’s good credentials, by the lights of the current proposal the inference is not truth preserving. Since both (4) and (6a) are beliefs expressible by critical assertions, the current proposal suggests

that their truth requires merely the existence of hallucinations, hypotheses or fictions with certain features. In particular, it requires only (i) that according to the relevant hallucination, hypothesis or fiction, Macbeth created Dag, and (ii) that according to the relevant hallucination, hypothesis or fiction, Dag is identical to Ger. At the same time, since (4e) only concerns ordinary real particulars and is *not* expressible by a critical assertion, its truth requires (iii) that Macbeth created Ger. Since (i) and (ii) can be true when (iii) is false, the inference at hand is not truth preserving. Thus, to accept the current proposal is to accept an intolerable rift between the validity and formal goodness of deductive inferences.

The third problem with the current proposal is that it makes it entirely mysterious how the truth conditions of our beliefs are determined. Think again of the beliefs that Macbeth created Dag and that Macbeth created Ger. Since there can be no single relation denoted by the concept “ x created y ” on which the truth of both beliefs turns, the concept must denote distinct relations in the two beliefs. So we must ask what determines which of the two relations the concept denotes in each case. But it appears that it is quite hard to give a plausible answer to this question. In fact, the only initially plausible answer would be that since the folk are willing to be committed to Ger’s existence, in believing that Macbeth created Ger they use the concept “ x created y ” with its “standard” denotation. But, since they are unwilling to be committed to Dag’s existence, in believing that Macbeth created Dag they use the concept “ x created y ” with an alternative denotation. This answer, however, would not satisfy the irrealist. As the examples of Ed and Irma from section 4.1.3 show, folk who believe that Macbeth created Dag might be quite willing to undertake a commitment to Dag’s existence. If so, in believing that Macbeth created Dag such folk would use the concept “ x created y ” with its “standard” denotation. Such folk would thereby believe something whose truth requires the existence of hallucinated entities, contrary to hermeneutic irrealism.

It seems then that we should not assign beliefs expressible by critical assertions truth conditions along the lines of (4a), (4b), (4c), (4d) or the appropriate instance of (R). Still, can such beliefs not be assigned some other truth conditions, in a way that is not vulnerable to the three objection just given? Recent work by Tim Crane (2013) and Mark Sainsbury (2010) suggests that there is a way to do this.⁶ In presenting this view, both thinkers have

⁶In previous sections Sainsbury’s work was read as recommending a form of paraphrase irrealism. Though

focused on assertions of the form

(8) x is famous.

Crane (2013, p. 135), and more explicitly Sainsbury (2010, p. 136-138), recommend analyzing statements of (8)'s form thus: Folk who assert an instance of (8) often do believe the instance's literal content. Thus, folk who assert "Dag is famous" often believe that Dag is famous. But the fact that the folk believe that Dag is famous does not mean that they are committed to Dag's existence. They are not so committed because the truth conditions of (8) are given by

(8a) For many people there is some propositional attitude A and some property F such that they A that x is F .

An instance of (8a) can be true when the substitution instance of " x " does not denote. Hence, folk who believe that Dag is famous believe something whose truth does not require Dag's existence.

If all critical assertions could be analyzed in the way instances of (8) just were, hermeneutic irrealism would have a version that is not vulnerable to the arguments of sections 4.1.3-4.1.5. But though ingenious, the suggested way of saving hermeneutic irrealism is untenable. It is untenable because ontologically "non-committal" truth conditions simply cannot be plausibly assigned to all critical assertions. In fact, they cannot even be plausibly assigned to instances of (8).

There are three reasons why (8a) cannot give the truth conditions of (8). First is the problem of non-specificity: Suppose that there is no particular dagger - whether merely hallucinated or real - that is widely known. But suppose also that Macbeth wants to hallucinate a dagger. His desire is non-specific, i.e., there is no particular dagger that he wants to hallucinate, and any old dagger would do for him. Finally, suppose Macbeth's desire

this seems to be the correct way to read him in general (assuming, as before, that one reads him as recommending hermeneutic irrealism at all), there is another reading of him that might be more suitable in some contexts. In his (2010), Sainsbury proposes various treatments of various assertions which appear to concern fictional entities. Some assertions are simply paraphrased, in the spirit of paraphrase irrealism. Others are paraphrased by the addition of an operator, in the spirit of operator irrealism. Others still are treated as true only relative to a presupposition, in the spirit of attitude irrealism. All these suggestions are vulnerable to variants of the arguments of sections 4.1.3-4.1.5. However, there is a final class of assertions that Sainsbury can be read as treating in new, fourth, way. It is this forth strategy that is now under consideration.

is well known, so that many people bear some attitude to the thought, “a dagger is what Macbeth wants to hallucinate”. If the truth conditions of (8) were given by (8a), it would follow from the last supposition that “a dagger is famous” is true. But this cannot be right. We are supposing that there is no particular dagger that is widely known; and if “a dagger is famous” is true, there must be a particular dagger that is widely known.

Second is the problem of substitution: Suppose Dag becomes famous after Macbeth goes on all the TV news shows and tells everybody, “I hallucinated a dagger, Dag”. Now let’s introduce a descriptive name (and equip ourselves with a descriptive singular concept) by stipulating that we will call whatever hallucinated dagger Macbeth in fact hallucinated in 1040 “*D*”. Since Macbeth’s TV performances made the dagger he hallucinated in 1040 well known, “*D* is famous” is true. But if so, the truth conditions of (8) are not given by (8a). The way the concept “*D*” was introduced guarantees that there are very few people that bear a propositional attitude to a thought in which “*D*” occurs. Therefore, by the lights of (8a), “*D* is famous” is false.

Third and last is the problem of existence: Suppose that though Macbeth never hallucinated any cupcakes, everybody erroneously believes that there is exactly one cupcake that he did hallucinate. In fact, suppose everybody believes that the cupcake Macbeth hallucinated is, in the hallucination, yummy. If the truth conditions of (8) were given by (8a), it would follow from the last supposition that “the cupcake Macbeth hallucinated is famous” is true. But this cannot be right. Since Macbeth never hallucinated any cupcakes, nothing satisfies the description “the cupcake Macbeth hallucinated”. And this fact is the source of our strong intuitions that “the cupcake Macbeth hallucinated is famous” is untrue.

These considerations illustrate why hermeneutic irrealism cannot be saved by assigning ontologically “non-committal” truth conditions to beliefs expressible by critical assertions.

4.2 A REALIST PROPOSAL

4.2.1 Abstract Realism

The arguments of section 4.1 show that hermeneutic irrealism is untenable. Critical assertions made by the folk express neither make beliefs nor beliefs with a content other than that of the critical assertions themselves. Instead, they express beliefs whose content is just the one literally asserted. Our intuitions that these assertions are true (simpliciter) likewise indicate that we believe what they say. And since critical assertions entail that hallucinated entities exist, it follows that we are committed to the existence of hallucinated entities, just as the descriptive part of realism suggests.

But if we are committed to the existence of hallucinated entities, we must have some conception of them, some idea of what these entities might be like. I will now argue that this conception is captured by abstract realism - a view inspired and promoted by the works of Kripke (2013), Salmon (1998), Schiffer (1996), Searle (1975), Thomasson (1999), and Van Inwagen (1977).⁷

To articulate our conception of hallucinated entities, abstract realists begin by noting that we can hallucinate particulars of two distinct kinds. First, we can hallucinate real particulars. For example, we can hallucinate things such as Barack Obama, or the Empire State building. Secondly, we can hallucinate particulars that are not real, i.e., that are *merely* hallucinated particulars. For example, Macbeth hallucinated Dag and Morgana. But to say that we can hallucinate both real and merely hallucinated particulars is not yet to articulate how we conceive of either group. To do that we need to revisit a second, orthogonal distinction, which we drew in section 4.1.1.

Section 4.1.1 distinguished between truth (simpliciter) and truth in a hallucination. Thus, for example, our intuitions are that:

- “Morgana created Dag” is (i) true in Macbeth’s hallucination, but (ii) false simpliciter;
- “Macbeth created Dag” is (i) false in Macbeth’s hallucination, but (ii) true simpliciter.

⁷Another well established realist position - concrete realism - was supported by Lewis (1978). Though I doubt that the view captures our intuitive conception of hallucinated entities, it is well motivated and deserves serious consideration. Since such consideration will, however, require a significant digression from the main thread, I will leave the view untouched for the remainder of this chapter.

Similarly, if I have a hallucination in which Barack Obama is Kenyan, then our intuitions say that:

- “Obama is Kenyan” is (i) true in my hallucination, but (ii) false simpliciter;
- “Obama is American” is (i) false in my hallucination, but (ii) true simpliciter.

The distinction between truth and truth in a hallucination is crucial to exploring our conception of hallucinated particulars. If our intuitions say that the assertion “ A is F ” is true (simpliciter), then our conception of A is one on which A has the property of being F . However, if our intuitions say that the assertion “ A is F ” is true in a hallucination h , then our conception of A is one on which A has the property of being F -in- h . And in general, having the property of being F is independent of having the property of being F -in- h . Thus, for example, our conception is that Obama has the property of being American, lacks the property of being Kenyan, and has the property of being Kenyan in my hallucination. Similarly, our conception is that Dag has the property of being created by Macbeth, lacks the property of being created by Morgana, and has the property of being created by Morgana in Macbeth’s hallucination.

If all this is on the right track, then the core of our intuitive conception of real and merely hallucinated particulars is to be derived from our intuitions about the truth (simpliciter) of assertions about these particulars. Here we can set real hallucinated particulars aside—our core conception of them is clear enough—and focus on merely hallucinated particulars. So consider Dag again. We already noted that, intuitively, Dag was created by Macbeth. Here are some other intuitions we have regarding Dag:

- “Macbeth hallucinates Dag” is (i) false in Macbeth’s hallucination (since the hallucination has Macbeth *perceiving* Dag rather than hallucinating it), but (ii) true simpliciter;
- “Macbeth created Dag by hallucinating it” is (i) false in Macbeth’s hallucination (since the hallucination has Morgana creating Dag), but (ii) true simpliciter;
- “Dag is metallic, sharp, and in front of Macbeth” is (i) true in Macbeth’s hallucination, but (ii) false simpliciter (since if Dag were metallic, sharp, and in front of Macbeth, anyone near Macbeth could prick Macbeth with it).

The mildest rational reconstruction of these intuitions seems to be the following: We conceive of Dag as an *abstract* and *mind-dependent* particular. It is because we conceive of Dag as abstract that we accept that we cannot perceive Dag and that we cannot grab it or prick Macbeth with it. Only concrete particulars are in the business of being perceptible and so

usable. And it is because we conceive of Dag as mind-dependent that we hold that Dag's creation could have been effected by Macbeth's hallucinating it.

What is true of Dag is true generally of merely hallucinated particulars. A mild rational reconstruction suggests we intuitively conceive of these particulars as abstract and mind-dependent particulars, created by those who hallucinate them first, on the occasion of their first hallucinating them. This conception of merely hallucinated particulars is very close to our conception of other abstract and mind-dependent entities, such as stories, passwords, patents, marriages, nations, monetary values and holidays. Our conception is that all these entities exist, but do so as abstract entities, contingently created by the mental activity of a person or a group. Merely hallucinated particulars are no different. Our conception of them is therefore not peculiar or novel. Rather, we take them to exist in the same manner that stories, passwords etc. exist - i.e., contingently, abstractly, and mind-dependently.⁸

Hallucinated particulars, however, do not exhaust the realm of hallucinated entities. We can also hallucinate properties and relations. For example, when he hallucinated Dag, Macbeth also hallucinated a certain dagger-like shape. Furthermore, we can hallucinate states of affairs (which, for brevity, I will hereafter refer to as "events").⁹ For example, when he hallucinated Dag, Macbeth also hallucinated Dag's being a dagger. So it remains to say how we conceive of hallucinated properties and events.

Regarding hallucinated properties and relations, it seems plausible to suppose that we can hallucinate most (if not all) the properties and relations we can perceive. For example,

⁸What is an abstract realist to say of particulars such as afterimages and phosphenes? On the one hand, these particulars do not fit comfortably into the mold of merely hallucinated particulars, since they can be experienced by someone even when her perceptual system is functioning normally, giving no indication of a hallucinatory state. On the other hand, these particulars are like merely hallucinated particulars, in being abstract and mind-dependent things, which exist contingently, in virtue of the fact that a person or persons created them in experiencing them. Given these two considerations, it seems best to simply say that when we experience after images or phosphenes, we are presented with certain existing mind-dependent entities. Whether we are presented with these entities via perceiving or hallucinating them is something we need not decide on. In fact, we might even say that we bear a third kind of experiential relation to them - experiencing*, say. For more on afterimages, see Ian Phillips' work (2012), which persuasively argues on empirical grounds that afterimage experiences are subjectively indistinguishable from perceptions of projected light phenomena.

⁹For present purposes, a state of affairs (an event) consists of the having of a property by a particular, or of several particulars bearing a relation to each other. An events is therefore an occurrence, i.e., that which occurs. Events can be denoted in various ways, very commonly by means of possessive-with-gerund constructions such as "*A*'s being *F*", or "*A*'s *R*-ing of *B*". Since events are occurrences, and not representations of or thoughts about occurrences, they are non-thinkables: One can only think *that A is F*. One cannot think *A's being F*.

we can both perceive and hallucinate (the property identical to) the shape of the Empire State building. The converse may not hold. Hurvich (1982), and following him Johnston (2004), have maintained that there is at least one property - supersaturated red - that we cannot perceive but can experience with the help of an afterimage. To quote Johnston (2004, p. 141-142):

There is a state that a subject can get into by being exposed to bright monochromatic unique green light (500 nanometers in wavelength) in an otherwise dark room for about twenty minutes. If we then turn the stimulus off, illuminate the room, and have the subject look at a small, not-too-bright achromatic surface, he will see a red afterimage. If the subject turns so that the afterimage is then superimposed on a small red background then something wonderful happens. The subject will then be afterimagining a supersaturated red, a red more saturated than any surface red one can see, a red purer than the purest spectral red light, light with a frequency of around 650 nanometers.

If this is true, then there may be properties that we can hallucinate but not perceive. Be that as it may, it seems plausible to suppose that our conception of properties which we hallucinate is no different from our conception of properties in general. So the fact that we can hallucinate properties poses no special difficulties that our conception of properties in general is not subject to.

Hallucinated events require closer scrutiny, since talk of hallucinated events can be quite confusing. Consider, for example, the assertion “Macbeth hallucinates Dag’s being sharp”. On the one hand, the assertion seems to be true (simpliciter), which suggests that Macbeth hallucinated some event denoted by “Dag’s being sharp”. On the other hand, we just said that Dag is an abstract particular, and therefore cannot be sharp. So we are left wondering - how can Macbeth hallucinate an event (Dag’s being sharp), which in principle cannot occur? To elucidate this matter, and with it our intuitive conception of hallucinated events, we must take a look at two ways we have of speaking about hallucinations - a literal way, and a hallucination-relative way.

Consider two contexts in which Macbeth might assert “Dag is sharp”. In the first, Macbeth says “I hallucinated a dagger into existence yesterday. I’ve named it Dag. Dag is a perfectly ordinary dagger. In particular, Dag is sharp”. If Macbeth were to assert this, we might question his sanity. The assertion indicates that Macbeth confuses a merely hallucinated entity with a real one, and literally says of it that it is sharp. Literally saying

that Dag is sharp is surely saying something false (simpliciter). Now consider a second context. In it, Macbeth says “I am hallucinating a dagger, Dag. Let me tell you how things are with Dag in my hallucination. Dag is sharp. ...” In asserting this, Macbeth is not confusing a merely hallucinated entity with a real one. His saying “Dag is sharp” is supposed to be understood as a specification of how the hallucination portrays Dag, not as a specification of how Dag really is. In other words, his assertion “Dag is sharp” is to be understood not literally, but in a hallucination relative way, as if it were implicitly prefixed with an “in my hallucination” operator. So in this second context, what Macbeth said in uttering “Dag is sharp” has the same truth conditions as a literal assertion of “in Macbeth’s hallucination, Dag is sharp”. Since the second assertion is true (simpliciter), so is what Macbeth said in the second context.

These considerations reveal that when engaged in discourse about hallucinations, we can use predicates not in a literal, but in an extended and hallucination relative way. Using a predicate in this hallucination relative way is using it to denote the in-the-relevant-hallucination counterpart of the property it literally denotes. In other words, uttering “ A is F ” in the course of hallucination discourse can express two kinds of predication. It can express a literal (henceforth: L) predication, which attributes being F to A . Alternatively, it can express a hallucination relative (henceforth: HR) predication, which attributes being F -in-the-contextually-relevant-hallucination to A .¹⁰ For example, an utterance of “Dag is sharp” can be taken to express either the thought “Dag is sharp_L” or the thought “Dag is sharp_{HR}”. The first is intuitively false (simpliciter), since it is intuitive that Dag is abstract and hence not sharp. The second is intuitively true (simpliciter), since it is intuitive that Dag has the property of being sharp in the relevant hallucination of Macbeth’s.¹¹

¹⁰A similar view has been suggested by Kripke (2013) and Salmon (1998). They suggest that when engaged in discourse about fictions, we often utter “ A is F ” to say that A is, in the relevant fiction, F .

¹¹Brock (2002, p. 12-13) and Kroon (1994, p. 213) have remarked that certain “mixed” assertions might pose a problem for views such as the one now being presented. The problem they raise can be illustrated by the utterance, “Dag is sharper than any real dagger”. If all real daggers are dull, and the utterance is made in the appropriate context, it would be natural to take the utterance to express a truth. But it is hard to see how it can express a truth. If it expresses the thought “Dag is sharper_L than any real dagger”, then it expresses an intuitively false thought. An abstract particular cannot be sharper than real daggers. If it expresses the thought “Dag is sharper_{HR} than any real dagger”, then it also expresses an intuitively false thought. By the semantics for HR predications, the last thought is true iff Macbeth hallucinates Dag as sharper than any real dagger. To hallucinate Dag in that way, Macbeth has to hallucinate not only Dag, but also all real daggers. And the trouble is that Macbeth hallucinates no real daggers. This difficulty

It will not have escaped the reader that L- and HR-predications bear systematic relations to truth (simpliciter) and to in-the-hallucination-truth: The thought “ A is F_{HR} ” is true (simpliciter) iff the thought “ A is F_L ” is true in the relevant hallucination. Similarly, “ A is F_{HR} ” is true (simpliciter) iff “ A is F -in-the-relevant-hallucination $_L$ ” is true (simpliciter). This fact makes HR-predications particularly useful, as we can use them together with L-predications to articulate thoughts that are intuitively true (simpliciter) in a non cumbersome way. For example, upon hallucinating Dag, Macbeth may utter, “a hallucinated dagger is before me”. This utterance intuitively expresses a truth. But it can only do so if it expresses the thought “a hallucinated $_L$ dagger $_{HR}$ is before $_{HR}$ me”, which is true (simpliciter) iff some hallucinated particular is (i) a dagger in the relevant hallucination, and (ii) before Macbeth in the relevant hallucination.

Let us now return to the topic of hallucinated event reports. Putting together all we have just noted, locutions of the form “ X hallucinates A ’s being F ” can be easily explained. Consider the utterance “Macbeth hallucinates Dag’s being sharp”, which can be understood as true (simpliciter). When so understood, “Dag’s being sharp” is taken to denote an event—namely, the event of Dag’s being sharp *in Macbeth’s hallucination*. There is no tension between Dag’s abstractness and the occurrence of *this* event. We can, after all, hallucinate abstract items in all sorts of ways. At the same time, this event can serve as the denotation of “Dag’s being sharp” if the utterance of “Dag’s being sharp” involves an HR predication. Such a predication is plausibly involved because “Dag’s being sharp” is preceded with the utterance of “Macbeth hallucinates...”, which makes Macbeth’s hallucination contextually relevant. In this setting, “Dag’s being sharp” is naturally taken to express the concept “Dag’s being sharp $_{HR}$ ”, which denotes Dag’s being sharp in Macbeth’s hallucination. All

can be resolved by treating “mixed” assertions in a manner analogous to the one we used for the assertion “Dag is sharp”. We note two points. First, facts concerning what is *fictionally* the case can supervene on a combination of facts concerning our hallucinations and facts concerning other things. For example, the fact that Dag is, in the relevant fiction, sharper than any real dagger can supervene on the fact that Macbeth hallucinated Dag as sharp and on the fact that all real daggers are dull. Second, we can sometimes report on facts concerning the contextually relevant fiction by using predicates in an extended and fiction relative (“FR”) way. To do so is to utter “ A is F ” to express the thought “ A is F_{FR} ”, which is true iff A has the property of being F -in-the-contextually-relevant-fiction. For example, we can sometimes utter “Dag is sharper than any real dagger” to express a thought which is true iff Dag is, in the relevant fiction, sharper than any real dagger. By putting these two points together we can see how it could be natural to take “Dag is sharper than any real dagger” to express a truth.

this notwithstanding, it is also possible (though somewhat difficult) to understand “Macbeth hallucinates Dag’s being sharp” as untrue. Such an understanding involves taking “Dag’s being sharp” to express the concept “Dag’s being sharp_L”, which is intuitively denotationless.

Analogous points apply to talk concerning hallucinated events which involve real particulars. Suppose I utter “I hallucinate Prince Charles’ and Lady Diana’s kissing of each other”. This utterance can be true in two ways. First, it can express the thought “I hallucinate Prince Charles’ and Lady Diana’s kissing_{HR} of each other”. This thought is true iff I hallucinate Charles and Diana as kissing each other. I could have such a hallucination even if Charles and Diana had never met. Second, the utterance can express the thought “I hallucinate Prince Charles’ and Lady Diana’s kissing_L of each other”. This thought is true iff I hallucinate one of the many events in which Charles and Diana really kissed each other. In hallucinating such an event I would be hallucinating an event whose occurrence is independent of anyone’s hallucinating it.

As a final example, suppose I utter (upon hallucinating Obama as Kenyan) “I am hallucinating Obama’s being Kenyan”. Despite the fact that Obama is American, it is possible to take me to be saying something true (simpliciter). So doing requires taking my utterance to express the thought “I am hallucinating Obama’s being Kenyan_{HR}”, which is true iff I am hallucinating the event of Obama’s being, in my hallucination, Kenyan. My utterance can also be taken (with some difficulty) to express an untruth. So taking it requires interpreting “Obama’s being Kenyan” as “Obama’s being Kenyan_L”, which (since Obama is not Kenyan) is denotationless.

We are now ready to spell out our intuitive conception of hallucinated events. Like hallucinated particulars, hallucinated events are of two kinds. We can hallucinate hallucination-independent events - as when I hallucinate Prince Charles’ and Lady Diana’s kissing_L of each other. Alternatively, we can hallucinate hallucination-dependent events - as when Macbeth hallucinates Dag’s being sharp_{HR} or when I hallucinate Obama’s being Kenyan_{HR}. The metaphysics of hallucination-independent events raises no special difficulties. The metaphysics of hallucination-dependent events can be elucidated thus: Suppose X hallucinates A ’s being F_{HR} . In hallucinating this, X does not, of course, make it the case that A has property F . Rather, in hallucinating A ’s being F_{HR} , X brings to existence (or, to occur-

rence) a contingent and mind-dependent event - namely, A 's being F -in- X 's-hallucination. This event is mind-dependent since its existence (or, occurrence) requires X 's being in the mental state of hallucinating A 's being F_{HR} . The event is contingent since X could have failed to be in that mental state.

4.2.2 The Nonexistence Objection

Section 4.2.1 just presented a mild rational reconstruction of the way we (the folk) think and speak of hallucinated entities. Call this intuitive conception of ours “abstract realism”.

An immediate difficulty for this realist reconstruction is the nonexistence objection. The objection is quite straightforward: Assertions like “Dag does not exist” are intuitively true. And since they are intuitively true, it is not part of our intuitive outlook that there are such things as hallucinated entities. So the descriptive half of the realist position—the view that we believe hallucinated entities exist—is false.

The nonexistence objection can be used to motivate views on which both “Dag does not exist” and “Dag is a merely hallucinated entity” are straightforwardly true. Here I have in mind intentional inexistence (or “IE”) views, such as those put forward by Alexius Meinong (1960), Terence Parsons (1980), Graham Priest (2005) and A. D. Smith (2002). Though different in many ways, all these views suggest the following: (i) Some particulars do not exist; or equivalently, some particulars are not. (ii) Nevertheless, we can successfully represent, refer to, and quantify over, the non-existing. (iii) When we hallucinate a particular that is not real, that merely hallucinated particular does not exist.

IE views straddle the realist-irrealist chasm. On the one hand, like realist views, IE views assign referents to terms like “Dag”. IE views therefore have the characteristically realist advantage of also assigning straightforward, disquotational, truth conditions to at least some assertions that concern merely hallucinated particulars. And since assertions with disquotational truth conditions are immune to the problems raised in sections 4.1.3-4.1.5, IE views are less vulnerable to refutation than hermeneutic irrealist views. On the other hand, like irrealist views, IE views are not committed to the existence of merely hallucinated entities. IE views therefore have the characteristically irrealist advantage of a sparse ontology. Finally,

unlike both realist and irrealist views, IE views can assign straightforward, disquotational, truth conditions to nonexistence claims. Realists struggle with “Dag does not exist” since they hold that Dag does exist. Irrealists struggle with “Dag does not exist” since they hold that “Dag” does not refer, which complicates accounting for the statement’s truth. IE views are free from both problems, since they say both that “Dag” refers, and that Dag has the property of not existing.

I will now argue that the nonexistence problem should neither drive us away from realism, nor drive us towards IE views. I will argue these points in reverse order.

Regarding IE views, I will only (mis)quote Russell to say that they have all the advantages of theft over honest toil. Consider the properties of being created by Macbeth, or being distinct from this dagger (pointing to a real dagger). If something has one of these properties, if it is one of these ways, then it is. In other words, whatever has one of these properties must also exist. To deny this is to hold a flatly unintelligible position. Now, intuitively, Dag has both these properties. Intuitively, Dag was created by Macbeth, and it is distinct from all real daggers. So intuitively, Dag exists. As a consequence, IE views must be rejected - at least insofar as they purport to capture our intuitive outlook.¹²

I now turn to the nonexistence problem itself, and to its force against realism. Here it must be granted that we do often take utterances like “Dag does not exist” to express true thoughts. But it is also true that there are contexts in which we take utterances like “Dag does not exist” to express false thoughts. Suppose two psychologists are discussing Macbeth’s history of hallucinations. One of them says, “Dag does not exist”. It is perfectly natural for the other to object: “That’s false. The hallucinated dagger, Dag, came into existence in 1040, when Macbeth was working through his murderous impulses.” Or, the other could object by saying: “There certainly is a such a hallucinated entity as Dag! Macbeth started hallucinating it shortly before his ascension to the Scottish throne.”

These considerations reveal that those who make the nonexistence objection to realism are only considering half of the relevant linguistic data. Once we consider things in full, we find that the target utterances are accepted as true in some contexts, but rejected as false in others. This finding, however, should not be taken to imply that our intuitions regarding

¹²For other criticisms of IE views, see Sainsbury (2010, chapter 3) and Fine (1984).

the existence of merely hallucinated entities are inconsistent. A far more plausible lesson to draw is that in different contexts we use the same utterances to express distinct thoughts. In other words, the target utterances are ambiguous, or polysemous.

There are a number of realist-friendly accounts of how the target utterances are ambiguous [see, e.g., Kripke (2013), Salmon (1998), Thomasson (2010)]. But since a full discussion of the topic is beyond the scope of this chapter, I will make do with sketching out my favorite account. It is a descendant of the account recommended by Thomasson in her (1999).¹³ On this account, in any context c , the verb “to exist” can be used to denote either the property of being, or whatever property is denoted in c by the verb “to be real”.¹⁴ Thus, “Dag exists” sometimes means that Dag, in the widest sense, is (equivalently, that Dag exists₁); and sometimes it means what “Dag is real” (equivalently, “Dag exists₂”) would mean in the context of utterance. Now according to realists, since Dag does exist₁, “Dag exists₁” is true in any context. But since in most contexts Dag does not count as real, “Dag exists₂” is false in most contexts.¹⁵

A virtue of this account is that it explains why (i) “Dag does not exist, and is merely a mind-dependent object” can sound true (simpliciter), whereas (ii) “Dag does not exist, and is in my hand” cannot. To see why this should be surprising, consider that being a mind-dependent object, just as much as being in my hand, intuitively implies existing₁. So suppose

¹³Thomasson later abandoned her (1999) account, but her reasons are irrelevant to the current project. For more on this, see her (2010).

¹⁴I am supposing that “to be real” is itself context sensitive. In some contexts, numbers count as real but fictional characters do not; in other contexts, the reverse is the case.

¹⁵Everett (2013, p. 150) has raised the following objection: The utterance “Bush exists but Raskolnikov doesn’t” means the same as “Bush exists but Raskolnikov doesn’t exist”. Therefore, “exists” can mean the same when predicationally applied to both “Bush” and “Raskolnikov”. Therefore, “exists” is not ambiguous. The objection, however, fails for two reasons. First, that “exists” can have the same meaning when predicationally applied to both “Bush” and “Raskolnikov” does not entail that “exists” is not ambiguous. It is perfectly possible for “exists” to be ambiguous, as long as it can take one (or both) of its meanings when predicationally applied to both “Bush” and “Raskolnikov”. And, in fact, this seems to be the case. An utterance of “Bush exists but Raskolnikov doesn’t exist” will typically be made in a context in which one has distinguished between real and merely fictional entities. So it is natural to interpret the utterance as meaning the same as “Bush exists₂ but Raskolnikov doesn’t exist₂”, or equivalently, “Bush is real but Raskolnikov isn’t real”. Second, since the two suggested meanings of “exists” are easily conflated, we should be reluctant to trust our intuition that “Bush exists but Raskolnikov doesn’t” must mean the same as “Bush exists₂ but Raskolnikov doesn’t exist₂”. I see no reason why there should not be *some* contexts in which the second utterance would mean something that the first cannot, namely, that Bush exists₁ but Raskolnikov doesn’t exist₂. Finally, Everett’s objection is particularly puzzling, as he himself (2013, p. 63-74) offers a semantic theory suggesting that “exist” sometimes can and sometimes cannot be used interchangeably with “real” and in reference to perfectly ordinary particulars.

that “to exist” could only be used to mean to exist₁. In that case, “Dag is merely a mind-dependent object” and “Dag is in my hand” would both imply something that contradicts “Dag does not exist”, making both (i) and (ii) sound untrue. Since this is not the case for (i), “to exist” must have a second meaning. Now consider what happens if we interpret (i) and (ii) by taking the relevant occurrences of “exist” to mean exist₂. In that case, (i) would straightforwardly sound true (simpliciter), since Dag is intuitively both not real and merely a mind-dependent object. At the same time, (ii) would continue to sound untrue. After all, it is intuitively impossible for something to be both in my hand and not real. So we have a simple account of why (i), but not (ii), can sound true (simpliciter).

These considerations show that the nonexistence objection should neither drive us away from realism, nor drive us towards IE views. Evidence suggesting that “to exist” is ambiguous supports realist-friendly replies to the objection, while internal tensions in IE views show that they cannot be accepted.

4.3 FOR ABSTRACT REALISM

4.3.1 A Default Theory

While sections 4.1 and 4.2.2 argued that non-realist views violate our intuitions regarding hallucinated entities, section 4.2.1 presented a kind of realism—abstract realism—that carefully tracks our intuitions regarding hallucinated entities. Put together, these sections reveal that abstract realism is the default view we, the folk, hold.

Of course, that abstract realism is the default view does not in itself show that we should continue to subscribe to it. Perhaps revisionary irrealists are right, and there are strong reasons to deny that there are such things as hallucination-dependent entities. So it is to the evaluation of this possibility that we must turn next.

Still, that abstract realism is the default view does determine who has the burden of proof. It is the revisionary irrealist, who is asking us to overthrow our intuitive conception, that must shoulder this burden. The abstract realist only has to plausibly defend our intuitions

against the revisionary arguments.

Despite the last dialectic observation, I will devote the rest of this section (4.3) to three positive considerations in favor of abstract realism. The first two concern phenomenological and epistemological issues in the philosophy of mind. The third has more to do with the philosophy of language. Considerations against abstract realism will be answered in section 4.4.

4.3.2 The Phenomenology of Hallucinations

Irrealists and abstract realists take very different lines about the phenomenology of hallucinations.

On the abstract realist line, hallucinated particulars exist. So it is natural for abstract realists to hold that when a subject (who is neither blindsighted nor otherwise impaired) hallucinates a particular, there is a way that the particular feels_L to the subject. If the subject auditorily hallucinates the particular, it sounds_L some way to the subject. If the subject gustatorily hallucinates the particular, it tastes_L some way to the subject. And if the subject visually hallucinates the particular, it looks_L some way to the subject. Thus, by abstract realist lights, certain abstract mind-dependent items can—much like William Alston’s (1999) mental images or Anil Gupta’s (2012) sense-images—play a role in the subjective life of a subject. For example, when Macbeth visually hallucinated Dag, there was a way that Dag looked_L to him, viz., sharp and bloody.

Contrast this with the irrealist line. Since irrealists hold that there are no merely hallucinated particulars, they must hold that no such particulars ever feel_L any way to any subject. By their lights, when Macbeth visually hallucinated Dag, there was no particular at all that looked_L sharp and bloody to Macbeth. Still, in having his hallucination, there was a way that things in general looked_L to Macbeth. Things looked_L to Macbeth exactly as they would were Macbeth perceiving a real dagger that was sharp and bloody. So though things looked_L a certain way to Macbeth, no particular thing looked_L to Macbeth a way that would account for this.

When these two lines of thought are thus compared, the abstract realist one emerges as

clearly superior. To see this, consider the following commonplace: In perceiving, things in general feel_L as they do to the perceiver partly because of how the perceived items feel_L to the perceiver. For example, if I were to see a sharp and bloody dagger, things would look_L to me a certain way partly because the dagger I would see would look_L sharp and bloody to me. So it should be clear that there are connections between the way things in general feel_L to a subject in perceiving, and the ways that the perceived items feel_L to the subject. Abstract realism merely extends these connections to the realm of hallucinations. According to realists, in hallucinating, things feel_L as they do to the hallucinator in part because of how the hallucinated items feel_L to the hallucinator. For example, in hallucinating, things looked_L as they did to Macbeth in part because there was a particular, Dag, that looked_L sharp and bloody to him. The abstract realist line on the phenomenology of hallucinations therefore simply generalizes on everyman's stance on the phenomenology of perception. This makes the realist line both simple and intuitive.

Things are different on the irrealist line of thought. Irrealists hold that the way things looked_L to Macbeth in his hallucinating is in no way connected with some particular's having looked_L sharp and bloody to him. Consequently, they must deny that in hallucinating, things feel_L as they do to the hallucinator in part because of how the hallucinated items feel_L to the hallucinator. And this last point exerts pressure over their position. If they choose to say that the commonplace with which the previous paragraph began is true, they must resist the generalization that the abstract irrealists are recommending. As a result, irrealists will have to offer distinct accounts of the phenomenologies of hallucinations and perceptions. But this is both complicated, and phenomenologically unmotivated. Alternatively, irrealists may choose to deny the commonplace. They may say that in both perceiving and hallucinating, the way things in general feel_L to one is in no part due to the way certain items feel_L to one. But to deny the commonplace is to take a seriously counterintuitive position.

A further consideration in favor of the abstract realist line is this: Even if it is not usually the case, hallucinations can [*pace* Fish (2009, chapter 4)] be phenomenally identical to perceptions. So suppose I am visually perceiving a sharp and bloody dagger, and suppose further that Macbeth's hallucination is phenomenally identical to my experience. Now to me, it is phenomenally as if some particular looks_L sharp and bloody. And as Macbeth had

a phenomenally identical hallucination, to him too it was phenomenally as if some particular looked_L sharp and bloody. But irrealists and realists differ in how well they can accommodate this fact about Macbeth.

The abstract realist has no difficulty with Macbeth's mental state. To Macbeth, it was phenomenally as if some particular looked_L sharp and bloody; and indeed, by realist lights, some particular - Dag - did look_L sharp and bloody to Macbeth. So at least this aspect of the phenomenology of Macbeth's hallucination—i.e., the phenomenology associated with the particulars that feel_L certain ways to us, as opposed to the ways that those particulars do feel_L to us—was not misleading. The irrealist, however, cannot say the same. The irrealist must say that to Macbeth, it was phenomenally as if some particular looked_L sharp and bloody, though no particular so looked_L to him. So by irrealist lights, to Macbeth, it was phenomenally as if what did not exist, did. Hence, the irrealist must accept a misleading aspect of the phenomenology of hallucination that the realist need not.

A final phenomenological consideration in favor of realism has to do with hallucinations of real particulars. Recall the case in which I hallucinate Obama. Since Obama definitely exists, there seems to be little reason to deny that when I hallucinate Obama, Obama looks_L to me a certain way. But if that is right, the irrealist position is again under pressure. The irrealist ends up saying, after all, both that when Macbeth hallucinated Dag, no particular looked_L the relevant way to Macbeth; and that when I hallucinate Obama, some particular does look_L the relevant way to me. Thus, the irrealist is forced to distinguish between the phenomenology of hallucinations of real particulars (wherein real particulars are felt_L) and the phenomenology of hallucinations of merely hallucinated particulars (wherein no such particulars are felt_L). This distinction seems to have no phenomenological merit. It is motivated by mere ontological prejudice.

4.3.3 The Epistemology of Hallucinations

Hallucinations often give rise to mental states that fail to be true. Suppose, for example, that I hallucinate an apple in front of me (where there in fact is no apple there). I might be taken in by this hallucination, and form the untrue (mock) belief, "that's an apple". This (mock)

belief would have a demonstratively referring element (“that”), through which I would be attempting to demonstratively refer to an apple - as I would do were I perceiving one. The (mock) belief would also have an attributive element (“is an apple”), through which I would be attempting to classify the purportedly referred-to particular as an apple - again, just as I would do were I perceiving one. But since there is no apple to be referred to, my (mock) belief would be untrue. So much is uncontroversial, and accepted by both abstract realists and irrealists.

Realists and irrealist do disagree, however, on the precise reason for which my (mock) belief is untrue. According to irrealists, the (mock) belief is untrue because it contains a demonstrative element that fails to refer. According to abstract realists, the demonstrative element does refer to something. Though I tokened the demonstrative element while attempting to attend to an apple, I in fact tokened it while attending (unbeknownst to me) to a merely hallucinated abstract particular, which is only an apple in my hallucination. The demonstrative element refers to it. So the reason my belief is untrue is not that it involves failed reference. Rather, taken in by my hallucination, I classified the abstract particular as an apple (i.e., as an apple_L), and not as an apple-according-to-my-hallucination (i.e., as an apple_{HR}). In so misclassifying the abstract particular, I ended up with a false belief.

It appears then that we are at an impasse. Although they disagree about some crucial details, both abstract realists and irrealists can account for cases in which hallucinations give rise to mental states that fail to be true. Does this mean that epistemological considerations give no advantage to either abstract realism or irrealism? I think the answer is “no”. I will now present cases in which hallucinations intuitively give rise to knowledge, or at least to true beliefs. Abstract realists can provide a simple and unified theoretical account of such cases, which preserves our intuitions. Irrealists, on the other hand, cannot. So it will emerge that abstract realism has an epistemological advantage.

Consider Macbeth’s hallucination of Dag again. Suppose Macbeth is aware that he is hallucinating, and so is not taken in by this hallucination. Instead, he responds to his hallucination by sincerely asserting the following:

- (9) That is, in my hallucination, a dagger.
- (10) I created that.
- (11) That is not a real dagger.

- (12) That thing is distinct from Lady Macbeth.
- (13) That thing symbolizes murder to me.
- (14) That thing terrifies me.

Intuitively, assertions (9)-(14) are all true and express, if not knowledge, then true beliefs. Abstract realism can easily account for this intuition. By abstract realist lights, when Macbeth asserts (9)-(14) he expresses beliefs with precisely the content of (9)-(14). The demonstrative elements in those beliefs all refer to the abstract particular Dag, and the attributive elements in those beliefs all classify him correctly. So the truth of these beliefs is straightforward.

Given the aforementioned, abstract realism can offer this simple and unified treatment of experiential demonstrative reference and experience-based belief: Hallucinations are just like perceptions in that both may equally enable a subject the opportunity of demonstrative reference to experienced entities, and of conceptual classification of these entities in experience-based beliefs. Just as perceptions afford the perceiver awareness of the perceived entities, hallucinations afford the hallucinator awareness of the hallucinated entities. And in both cases, if the subject (i) attends to or experientially tracks the entities to which he is afforded awareness, and (ii) has the requisite conceptual capacities, she can demonstratively refer to these entities and conceptually classify them in various ways. Abstract realists can therefore hold that experiential demonstrative reference and the formation of experience-based belief is accomplished by a single set of attentive and conceptual capacities and proceeds in the same manner regardless of whether the subject is perceiving or hallucinating, and regardless of whether the subject is taken in by the hallucination or not.

Importantly, the abstract realist account just presented is consistent with various conceptions of experiential states. Here are two illustrations of this idea. First, representationalists hold that both perceptual and hallucinatory experiences represent things as being certain ways, and can be evaluated for veridicality at various (centered) possible worlds.¹⁶ An abstract realist who shares this view might say that awareness of perceived and hallucinated entities is afforded when these entities are represented in the content of the relevant ex-

¹⁶Representationalism is championed by Tyler Burge (1986; 2003a; 2010), Martin Davies (1992; 1997), Gareth Evans (1982), Gilbert Harman (1990), Christopher Hill (2009), John McDowell (1996), Colin McGinn (1991), Ruth Garrett Millikan (2004), Adam Pautz (2010), Christopher Peacocke (1992; 2001; 2004), James Pryor (2000), John Searle (1983), Susanna Siegel (2010) and Michael Tye (1995; 2000; 2009).

periences. Alternatively, the abstract realist might say that awareness of perceived and hallucinated entities is afforded when these entities are appropriately related to the content of the relevant experiences. Either way, the realist-representationalist can go on to say that this awareness can enable experiential attention, tracking, successful demonstrative reference and conceptual classification. So realism and representationalism are indeed consistent.¹⁷

Second, relationalists hold that perceptual experiences are presentational mental states, which immediately present perceived entities to the subject. Currently, most relationalists are also irrealists.¹⁸ But relationalism can also be conceived in combination with abstract realism. So conceived, relationalism says that hallucinations - just like perceptions - are presentational mental states. In other words, hallucinations immediately present hallucinated entities to the subject. A relationalism of this kind will suggest that when presented either with hallucinated or with perceived entities, the subject is afforded awareness of those entities. This awareness may in turn enable experiential attention, tracking and successful demonstrative reference. So realism and relationalism are consistent as well.

To summarize: Abstract realism suggests that the epistemology of hallucination-based belief is largely an extension of the epistemology of perception-based belief. Hallucinations - just like perceptions - afford the subject awareness of hallucinated entities. And like perceptual awareness, hallucinatory awareness can, under certain circumstances, enable the subject to attend to, track, demonstratively refer to, and conceptually classify the experienced entities.¹⁹

¹⁷Adam Pautz (2007) has raised some doubts regarding this point. I will address his concerns in section 4.4.

¹⁸Most relationalists accept metaphysical disjunctivism - a view on which perceptions do, but hallucinations do not, present and afford awareness of entities. Metaphysical disjunctivists therefore do not recognize merely hallucinated entities. Champions of disjunctivism are Bill Brewer (2011), John Campbell (in his and Quassim Cassam's 2014), William Fish (2009), and Michael G. F. Martin (2002; 2004). Matthew Conduct (2012) and Mark Johnston (2004; 2006) are also relationalists and irrealists, but they are not metaphysical disjunctivists in the full sense. On their view, hallucinations present and afford awareness of properties, but not of particulars. Their view is still an irrealist one, as Conduct and Johnston only recognize entities that exist independently of one's hallucinations. They do not, e.g., recognize merely hallucinated particulars. All this is not to say, however, that there are not relationalists that are also realists. William Alston (1999) and Anil Gupta (2012) in particular are relationalists who believe that in both perceptual and hallucinatory experiences we are immediately presented with existents. They further hold that in some hallucinatory experiences, at least some of the existents the hallucinator is presented with are not ordinary real ones. Instead, they are "mental images" (Alston) or "sense images" (Gupta). Alston and Gupta are therefore realists. It should be noted, though, that Alston and Gupta may differ substantially from each other and from *abstract* realists in the precise way that they conceive of these "mental images" or "sense images".

¹⁹It might be added, though, that there is the following difference between the epistemology of

Apart from offering a simple and unified epistemological theory, abstract realism also has the advantage of vindicating many of our epistemological intuitions: (i) The feelings that we can sometimes attend to, track and demonstratively refer to hallucinated entities are deemed accurate by this account. (ii) The feeling that we are often taken in by our hallucinations is likewise deemed accurate, as the account suggests that we often misclassify the entities we hallucinate. (iii) The feeling that we sometimes gain hallucination-based knowledge or true belief is deemed accurate, as the account makes room for our correctly classifying the entities we hallucinate. (iv) Finally, the account accomplishes all this while allowing the content of hallucination-based beliefs to be straightforwardly reflected by the semantic content of the assertions used to express those beliefs.

Now contrast the abstract realist epistemological outlook with the irrealist one. The immediate question for the irrealist is whether Macbeth's assertions (9)-(14) express knowledge or true beliefs. If the irrealist says "no", she is asking us to accept what is surely wrong. When asserting (9)-(14) Macbeth is surely giving expression to something which he gets correctly, not incorrectly, about how things stand. On the other hand, if the irrealist says "yes" she must specify what true beliefs Macbeth's assertions express. This might be difficult, since the irrealist must hold that the content of those beliefs is not straightforwardly reflected by the content of (9)-(14), which are (by her lights) untrue due to demonstrative reference failure.

How plausible is it then the the irrealist could account for the true beliefs which (9)-(14) express? Not at all, I feel. Three difficulties that stand in her way are these: (i) Assertions (10)-(14) are critical assertions (in the sense of section 4.1.1), and therefore, the arguments of sections 4.1.3-4.1.5 all apply to them. These arguments undermine the standard irrealist reinterpretations of such assertions. (ii) Macbeth could have replaced the demonstratives "that" or "that thing" in (9)-(14) with "Dag". This shows that Macbeth is expressing a thought concerning a hallucinated particular, rather than the hallucination itself or any other referential surrogate the irrealist might try to appeal to. (iii) Were Macbeth taken in by his hallucination of Dag and (mock) think "that's a (real) dagger", he would clearly be

hallucination-based belief and that of perception-based belief: In cases of hallucination, but not of perception, one is at great risk of being misled by one's experience into misclassifying the experienced entities.

attempting to think demonstratively about some particular, rather than about his mental state or any other referential surrogate. That is just one of the consequences of being taken in by a hallucination. But if Macbeth would be attempting to think demonstratively about a particular when taken in by a hallucination, why should he not attempt the same when he is not taken in by a hallucination? In other words, why should the thoughts Macbeth expresses in asserting (9)-(14) not equally count as ordinary demonstrative thoughts? If they do so count, however, their truth requires there to exist a demonstrated particular, contrary to irrealist contention.

Even if the irrealist could explain how (9)-(14) express true beliefs, there is still the worry that the irrealist position is neither as simple nor as intuitively accommodating as the abstract realist position. It is not intuitively accommodating because the irrealist cannot vindicate the feeling that we can attend, track and demonstratively refer to merely hallucinated particulars. By her lights, such particulars do not exist. The irrealist position is not simple because she denies that hallucinations afford the subject awareness of merely hallucinated particulars. In denying this, the irrealist forfeits the theoretically pleasing option of holding that the epistemology of hallucination-based belief is largely an extension of the epistemology of perception-based belief.

A final concern for the irrealist is best illustrated by reconsidering the case in which I hallucinate Obama. Suppose I know that Obama is not physically in front of me and I am only hallucinating him. Suppose further that given that this is the case, I respond to my hallucination by sincerely asserting to myself:

(15) That is Obama. I am hallucinating him.

Since Obama is a concrete mind-independent perceptible object with which I am familiar through pictures and other means, there seems to be little reason to deny that I can attend and demonstratively refer to him upon hallucinating him. So let us accept that the first sentence in (15) expresses my true demonstrative belief, of Obama, that he is Obama. If this is right, however, it follows that hallucinations can give rise to true demonstrative beliefs when those concern hallucinated *real* particulars. This consequence puts the irrealist at a disadvantage. The irrealist must deny, after all, that hallucinations can give rise to true

demonstrative beliefs which concern *merely* hallucinated particulars. (Such particulars don't exist, by irrealist lights.) So the irrealist is pushed to draw a distinction between the beliefs expressed by (9)-(14) and the belief expressed by the first sentence in (15). This distinction seems to have no epistemological merit. It is motivated by mere ontological prejudice.

4.3.4 The Semantics of Hallucinations

Consider the verb “to hallucinate”. This verb can be complemented both with noun phrases that pick out thinkables (e.g., “Macbeth hallucinates that Fido is barking”), and with noun phrases that denote non-thinkables (e.g., “Macbeth hallucinates Fido”). When “to hallucinate” is complemented with a phrase denoting a non-thinkable, the construction displays two interesting features. First is substitutional transparency: When x is a non-thinkable, “ A hallucinates x ” entails that A hallucinates y , if x and y are identical. For example, “Macbeth hallucinates Napoleon” entails that Macbeth hallucinates the person who was Emperor of the French from 1804 to 1815. Second is specificity: If F s are non-thinkables, “ A hallucinates an F ” entails that A hallucinates a particular F . In this way “to hallucinate” is unlike “to seek”. “Macbeth seeks an emperor” need not entail that he seeks any particular emperor. His seeking can be such that any emperor would do for him. But “Macbeth hallucinates an emperor_L” always entails that Macbeth hallucinates a particular emperor. The particular emperor Macbeth hallucinates is the one Macbeth might try to attend to, pick out demonstratively, or describe as “the emperor I am hallucinating here and now”. Similarly, “Macbeth hallucinates a dagger_{HR}” always entails that Macbeth hallucinates a specific thing (which, according his hallucination, is a dagger). As before, Macbeth might try to attend to this thing, pick it out demonstratively, or describe it as “the dagger-according-to-my-hallucination that I am hallucinating here and now”.²⁰

²⁰Graham Priest (2005, p. 67) also holds that “to hallucinate” has a use on which it is specific. Sainsbury, however, denies this. According to him (2010, p. 135), if I hallucinate a monster, it can be entirely appropriate for me to answer the question “which monster was it?” with “no particular monster. Just some horrible monster or other.” But though the response may be appropriate in some contexts, as a way of conveying that the hallucinated monster was not hallucinated as having any particularly remarkable features, it does not show that the intended use of “to hallucinate” can fail to be specific. If I hallucinate a monster, there is a particular one which seems to me to be perceived. And if I hallucinate a monster, I can give myself a true and definite answer to the question “which monster was it?” - the answer “that monster”. So, *pace* Sainsbury, there is no reason to think that where F s are non-thinkables, “ A hallucinates an F ” can fail to

Some semantic account of “to hallucinate” is needed if we are to explain why, when “to hallucinate” is complemented with a phrase denoting a non-thinkable, the resulting construction is specific and substitutionally transparent. An advantage of abstract realism is that it can provide a particularly straightforward semantic account for this kind.

By abstract realist lights, when “to hallucinate” is complemented with a phrase denoting a non-thinkable, the verb is used to denote a well-behaved relation, obtaining (at a world w , at a time t) between A and x iff A and x both exist at w at some time, and A hallucinates x at w and at t . This simple semantic account easily explains both the specificity and the substitutional transparency of the relevant constructions. Regarding specificity: When “hallucinates” denotes a relation between subjects and other entities, the truth of “ A hallucinates an F ” requires the relation to obtain between A and some entity, which is an F . But every entity is a specific and particular entity. So the truth of “ A hallucinates an F ” requires the relation to obtain between A and some particular F . Regarding substitutional transparency: When “hallucinates” denotes a well-behaved relation, the truth of “ A hallucinates x ” requires the relation to obtain between A and x . But if x and y are identical, then the relation also obtains between A and y .

A further advantage of the abstract realist semantic account of “to hallucinate” has to do with the verb’s existential transparency. According to abstract realism, when x is a non-thinkable, “ A hallucinates x ” entails that x exists (where “exists” is read tenselessly).²¹ For example, “Macbeth hallucinates Dag” entails that Dag exists. Existential transparency follows from the abstract realist semantic account, since the account says that when used in the relevant way, “to hallucinate” denotes a relation between entities that exist at some time. As a result, the truth of the relevant “ A perceives x ” constructions requires the (tenseless) existence of x . This consequence of the abstract realist semantic account is, *pace* irrealist protests, an intuitive one. As section 4.1 showed, the folk hold that hallucinated entities exist. The abstract realist semantic account therefore only vindicates their intuitions.

One way to summarize all this is to say that by abstract realist lights, the semantics of “to hallucinate” closely resembles the semantics of “to perceive”. Like “to hallucinate”, “to

be specific.

²¹Tenselessness is important. I can hallucinate Napoleon though he has by now ceased to exist.

perceive” can be complemented with noun phrases that pick out thinkables (e.g., “Macbeth perceives that Fido is barking”), and with noun phrases that denote non-thinkables (e.g., “Macbeth perceives Fido”). And like “to hallucinate”, when “to perceive” is complemented with a phrase denoting a non-thinkable, the resulting construction is existentially transparent, substitutionally transparent, and specific. In the case of “to perceive”, all would agree that these three features are explained by the fact that when appropriately complemented, the verb denotes a relation obtaining (at a world w , at a time t) between A and x iff A and x both exist at w at some time, and A perceives x at w and at t . Abstract realism simply extends this easily acceptable explanation to the case of “to hallucinate”.

Let us turn now to the irrealist side of things. Irrealists will of course deny that “to hallucinate” is existentially transparent. But they must still account for the verb’s specificity and substitutional transparency. It is unclear whether they can do that, and if they can, what elaborate resources they will have to appeal to. The matter is difficult to investigate further given the limited extent to which the semantics of “to hallucinate” has been studied. Still, one can appreciate the difficulties facing the irrealist by exploring what seems to be the most careful and thoughtful irrealist study of “to hallucinate” currently available. I have in mind Mark Johnston’s “The Obscure Object of Hallucination”.

Johnston tries to account for the phenomenon of visually hallucinating particulars thus (2004, p. 156-7):

As a first pass at an account of hallucinating particulars, we can say that

X’s visually hallucinating @

where ‘@’ holds a place for a designator which putatively designates a particular, consists in this:

There is some sensible profile of which X is visually aware,
And X is not visually aware of instantiations of this profile in the scene before his eyes,
And either this profile strikes X as being @ or X’s visual awareness of the profile is caused in the right way by an earlier perception of @ or by thought to the effect that @ is such and so.

Notice that the context ‘strikes X as being...’ is intensional, in the bland sense that it does not follow from the fact that something, say a dutiful father, strikes X as being, say, Santa Claus that Santa Claus exists.

Johnston's account would satisfy an irrealist, since it allows for "*A* hallucinates *x*" to be true even when "*x*" is denotationless. But the account is also severely limited, as it only provides truth conditions for assertions of the form "*A* visually hallucinates *x*" in which "*x*" putatively designates a particular. Assertions of the form "*A* visually hallucinates an *F*" are left untreated. As a result, it is impossible to test if Johnston's account can accommodate the specificity of "to hallucinate".

A more severe problem with Johnston's account is that it runs afoul of the substitutional transparency of "to hallucinate". Here are two illustrations of this: (i) Suppose Alice's thinking "Obama is an American" causes her (in the right way) to be visually aware of some appropriate sensible profile *P*, though she is not visually aware of instantiations of *P* in the scene before her eyes. By Johnston's account, it follows from this supposition that "Alice visually hallucinates Obama" is true. Now let's introduce a descriptive name (and equip ourselves with a descriptive singular concept) by stipulating that we will call whoever is the 44th president of the United States "*O*". This stipulation guarantees that *O* is Obama. So, in accordance with the substitutional transparency of "to hallucinate", we find that "Alice visually hallucinates *O*" is intuitively true. But Johnston's account cannot vindicate this intuition. The way the concept "*O*" was introduced guarantees that Alice has no thoughts in which "*O*" occurs, and that *P* does not strike Alice as *O*. Supposing also that Alice's visual awareness of *P* was not caused by an earlier perception of *O* (i.e., of Obama), Johnston's account entails that "Alice visually hallucinates *O*" is false. (ii) Suppose in 1050 Macbeth comes to be visually aware of some dagger-like sensible profile *P*^{*}, though he is not visually aware of instantiations of *P*^{*} in the scene before his eyes. Suppose further that since Macbeth retains a visual memory of his 1040 hallucination of Dag, *P*^{*} strikes him as being Dag. By Johnston's account, it follows from these suppositions that "Macbeth visually hallucinates Dag" is true. Now in a previous section we introduced the descriptive name (and equipped ourselves with the descriptive singular concept) "*D*" by stipulating that *D* is whatever hallucinated dagger Macbeth in fact hallucinated in 1040. This guarantees that *D* is Dag. So, in accordance with the substitutional transparency of "to hallucinate", we find that "Macbeth visually hallucinates *D*" is intuitively true. But as before, Johnston's account cannot vindicate this intuition. The way the concept "*D*" was introduced guarantees

that Macbeth has no thoughts in which “ D ” occurs, and that P^* does not strike Macbeth as D . D is furthermore not a perceptible object, and so Macbeth has no earlier perceptions of D . As a result of all this, Johnston’s account entails that “Macbeth visually hallucinates D ” is false.

I hope these considerations against Johnston’s account illustrate the difficulties facing irrealist treatments of the semantics of “to hallucinate”.

4.4 RESPONSES TO DOUBTS

Section 4.3 has argued for abstract realism by exposing its theoretical advantages for phenomenology, epistemology and semantic analysis. But some readers might not be swayed by these advantages, as long as certain irrealist doubts of theirs are left unanswered. I will therefore close this chapter with a section attempting to answer six of these doubts.

Doubt 1. Merely hallucinated particulars are metaphysically obnoxious, and out of step with the rest of physical nature. So they simply can’t exist. Besides, since these particulars are abstract and so causally inert, how can we know anything about them, or even refer to them?

Reply. Merely hallucinated particulars are not metaphysically obnoxious. It is true that they are abstract, mind-dependent and contingent existents. But we already noted that in this they are no different from stories, passwords, patents, marriages, nations, monetary values and holidays. As the metaphysical respectability of stories, passwords etc. is beyond question, so is that of merely hallucinated particulars.

Merely hallucinated particulars are furthermore not out of step with the rest of physical nature. Consider again the analogy with stories, passwords etc. Physicalists say that the facts about the existence and properties of stories, passwords etc. globally and metaphysically supervene on the physical facts.²² For this reason, they also say that stories, passwords etc. are neither primitive nor recalcitrant additions to our ontology, out of step

²²In other words, physicalists say that if w is a metaphysically possible world and a minimal physical duplicate of our world, then the facts about the existence and properties of stories, passwords etc. which obtain in our world also obtain in w .

with physical nature. The same is true of merely hallucinated particulars. By physicalist lights, facts about the existence and properties of merely hallucinated particulars globally and metaphysically supervene on the physical facts.²³ For example, by physicalist lights, the complete physical facts about the nature, causes and effects of Macbeth's hallucination metaphysically determine that Dag was created by Macbeth. Since facts about the merely hallucinated supervene on the physical, physicalists should be happy to admit that merely hallucinated particulars are neither primitive nor recalcitrant additions to our ontology.

It remains to address the concern that we cannot have knowledge about, or even refer to, causally inert merely hallucinated particulars. Here too the response appeals to the analogy with stories (or passwords ...). Stories are abstract and causally inert. Nonetheless, if I create a story and have the requisite skill and opportunity, I might name it and thus equip myself with a way of referring to it. Having created my story, I might also gain various kinds of knowledge about it. And I might communicate this name and knowledge to others, thus extending the circle of those who know about and can refer to my story. So knowledge about and reference to stories is not impossible. Rather, both are made possible in virtue of the fact that stories are created by people. The creator of a story is so related to the abstract particular, that it is no mystery that (under certain conditions) this creator can know about and refer to it.

What is true of stories is also true of merely abstract particulars. When Macbeth hallucinated Dag, he also named it. He thus equipped himself with a way of referring to it. In hallucinating it, Macbeth was also afforded an opportunity to gain knowledge about Dag. (E.g., that Dag is, in Macbeth's own hallucination, a dagger.) Thus, knowledge about and reference to merely hallucinated particulars is not impossible. Rather, both are made possible in virtue of the fact that such particulars are *hallucinated* by people. The hallucinator of a merely hallucinated particular is so related to the abstract entity, that it is no mystery that (under certain conditions) this hallucinator can know about and refer to it.

Doubt 2 [due to Crane (2013, chapter 5) and Sainsbury (2010, p. 132-143)]. Purported facts about the existence and properties of merely hallucinated particulars are not primitive.

²³In other words, physicalists say that if w is a metaphysically possible world and a minimal physical duplicate of our world, then the facts about the existence and properties of merely hallucinated particulars which obtain in our world also obtain in w .

Instead, they globally and metaphysically supervene on other facts. These other facts are the physical ones, or perhaps just those about the psychological states and episodes of subjects. Either way, it follows that we can eliminate every description of a purported fact about a merely hallucinated particular by replacing it with a description of the facts on which it supervenes. Does that not show that we can do away with merely hallucinated particulars after all?

Reply. Facts concerning stories (or passwords ...) globally and metaphysically supervene on other facts. (For example, that Edgar Allan Poe created *The Murders in the Rue Morgue* plausibly supervenes on the total psychological facts.) So perhaps we can eliminate every description of a purported fact about a story (password ...) by replacing it with a description of the facts on which it supervenes. But even if we can do this, it does not follow that stories (passwords ...) do not exist. (Neither does it follow that Poe did not create *The Murders in the Rue Morgue*.) Besides, there are good reasons to recognize the existence of stories (passwords ...), not the least of which is that these things are a part of commonsense folk ontology.

What is true for stories, passwords etc. is also true for merely hallucinated particulars. Even if we could eliminate talk of them in the envisioned way, that would not show that merely hallucinated particulars do not exist. Section 4.3 provided a few theoretical reasons to recognize the existence of merely hallucinated particulars. Another reason to recognize this, shown by sections 4.1 and 4.2.2, is that merely hallucinated particulars are a part of commonsense folk ontology.

Doubt 3 [due to Johnston (2004) and Conduct (2012)]. We can fully account for the phenomenal character of hallucinations by simply taking them to be states that afford us experiential awareness of complexes of sensible properties and relations, in which no merely hallucinated particulars take part. So we lose nothing if we say that there are no hallucinations of merely hallucinated particulars. This way, we can go on holding that merely hallucinated particulars don't exist. Isn't that a metaphysically more economic way of doing things?

Reply. Whether we can or cannot account for the phenomenal character of hallucinations in the suggested way is beside the point. The crucial problem with the suggestion is that it

ignores the epistemological fact that we can attend, track, demonstratively refer to, and gain knowledge of merely hallucinated particulars. The suggestion also ignores the transparent semantics of the verb “to hallucinate”. To account for these, merely hallucinated particulars are required to exist.

Doubt 4 [due to Johnston (2004)]. You can hallucinate the color Sienna Red even if you have never seen it, and even if no causal chains link you with it. And under certain conditions, a hallucination of Sienna Red can be the sole source of your ability to have de re thoughts about the color. Contrast this with hallucinating Obama. A person in no way causally linked with Obama (who, in particular, has never perceived Obama) cannot hallucinate him. And no hallucination had by such a person can be the sole source of an ability to have de re thoughts about Obama. Don't considerations such as these show that hallucinations can afford hallucinators experiential awareness of properties, but not of particulars?

Reply. It is indeed plausible that someone who is in no way causally linked with x —where x is an ordinary real particular (e.g., Obama)—can have no hallucination that counts as a hallucination of x . But similar constraints do not apply to merely hallucinated particulars. Macbeth, for example, was never causally linked with Dag, as Dag is causally inert. Nevertheless, Macbeth had a hallucination of Dag. So it appears that there are differences between real particulars on the one hand, and merely hallucinated particulars and properties like Sienna Red on the other hand, with regards to the conditions under which they may be hallucinated.

It is important, however, not to overstate differences between various kinds of hallucinated entities. In particular, the differences of the last paragraph only seem to concern the conditions under which one can hallucinate various entities. Once a subject hallucinates an entity, the kind of entity that it is seems to make no difference with respect to the subject's experiential awareness of, or ability to have a *de re* thought about, the entity. For example, hallucinations of Obama, Dag or Sienna Red equally afford hallucinators experiential awareness of—and (perhaps) the ability for *de re* thought about—Obama, Dag and Sienna Red. It is just that without the right causal connection, one cannot hallucinate Obama; though Macbeth could hallucinate Dag and Sienna Red.

Doubt 5 [due to Pautz (2007; 2010)]. Hallucinations can be contradictory and indetermi-

nate. For an example of a contradictory hallucination, suppose that you have a hallucination with the same phenomenal character as in the waterfall illusion, where you perceive an object which seems to you to be both moving and stationary. For an example of an indeterminate hallucination, suppose that you have a hallucination with the same phenomenal character as when you perceive a pink object in the periphery of your visual field, so that it seems to you to be pink, though no specific shade of pink. By realist lights, the possibility of such contradictory and indeterminate hallucinations requires there to be hallucinated particulars with indeterminate and contradictory properties. But such strange objects surely do not exist.

Reply. The abstract realist can answer this doubt along two distinct lines:

On the first line of response, the abstract realist concedes that in certain contradictory or indeterminate hallucinations, the subject is experientially aware of hallucinated particulars with some unusual properties. But the abstract realist insists that even in such cases, the particulars do not have indeterminate or contradictory properties. Hence, there is no reason to doubt their existence.

To see how this line of response might be elaborated, consider Pautz's two examples again. In the first example, what you hallucinate is not a particular that is both (i) moving, and (ii) stationary. Such a particular would be impossible. Rather, what you hallucinate is a particular, that is merely both (i) according-to-the-hallucination-moving, and (ii) according-to-the-hallucination-stationary. There is no impossibility associated with something's having both these properties. Particulars that have them merely reveal that, *in a hallucination*, things can be impossible ways. In the second example, what you hallucinate is not a particular that is both (i) pink, and (ii) no specific shade of pink. Such a particular would be strangely indeterminate. Rather, what you hallucinate is a particular, that is merely both (i) according-to-the-hallucination-pink, and (ii) not according-to-the-hallucination- S (for any specific shade of pink, S). Again, there is no impossibility associated with having both these properties. Particulars that have them merely reveal that, *in a hallucination*, things can be indeterminate ways.

On the second line of response, the abstract realist does not even concede that, in a hallucination, things can be impossible or indeterminate ways. Rather, the realist insists that in any hallucination h , the subject is experientially aware of particulars that are, in h , perfectly

possible and determinate ways. So why do things sometimes seem to be contradictory or indeterminate ways to hallucinating subjects? Because our cognitive systems may react to certain ordinary episodes of (hallucinatory) experiential awareness, by making things seem to us to be contradictory or indeterminate ways.

Pautz's two examples will again help show how this line of response might be elaborated. Start with the first example, which calls for a brief look at the waterfall illusion. The waterfall illusion is brought about by first fixating on a moving particular for a while, and then looking at a stationary one. Subjects who do this will both perceive the stationary object's being stationary, and experience an after-image's being, in the experience, in motion. Each of these two experiences is unremarkable in itself. In particular, experiencing an after-image's being, in the experience, in motion is a way of being aware of something's being, in the experience, a perfectly possible and determinate way. Nonetheless, to subjects who simultaneously have both experiences, the stationary object seems to be both stationary and moving. This shows that things' seeming to you to be an impossible way can result not from your experiential awareness of anything unusual, but from the way your cognitive system responds to such awareness.

So much for the waterfall illusion. Now consider Pautz's example of a hallucination with the same phenomenal character as the waterfall illusion. Since it matches the waterfall illusion in character, it is plausible that such a hallucination involves the subject in two episodes of experiential awareness. In the first, the subject is aware of some hallucinated particular's being, in the hallucination, stationary. In the second, the subject is aware of a distinct item - a hallucinated after-image. More specifically, the subject is aware of the after-image's being, in the hallucination, in motion. Thus, in each of the two episodes, the subject is aware of an item's being, in the relevant hallucination, a perfectly possible and determinate way. Still, in simultaneously undergoing both episodes, the first hallucinated particular seems to the subject to be both stationary and moving. It seems to the subject to be this impossible way as a result of the operation of the subject's cognitive system. But the subject is not experientially aware of any one thing's being, in the hallucination, an impossible way.

Consider Pautz's second example next. In this example, the subject has a hallucination

with the same phenomenal character as when you perceive a pink object in the periphery of your visual field, so that it seems to you to be pink, though no specific shade of pink. Here the abstract realist might say that despite appearances, in undergoing such a hallucination the subject is made experientially aware of some hallucinated particular's being, in the hallucination, a specific shade of pink - Orchid Pink, say. It is only that, since the particular is hallucinated as in the periphery of the subject's visual field, the particular does not seem to the subject to be Orchid Pink, but merely pink. So construed, the example is one in which the subject hallucinates things as being (in the hallucination) a perfectly determinate way. Things seem to the subject to be an indeterminate way only because the subject's cognitive system fails to fully felicitously respond to the hallucination.

Inspired by Pautz (2007, p. 511), one might raise the following three objections: First, it is implausible that anything that happens in the brain determines that a hallucinated particular is, in the hallucination, Orchid Pink rather than Champagne Pink. Second, if the hallucinated particular is, in the hallucination, Orchid Pink, why can the subject not tell that it is? Third, the suggestion requires the supposition that the phenomenal character of an experience is not entirely determined by what the subject is experientially aware of, but is also influenced by the way things seem to the subject to be.

Regarding the first objection: It is an empirical matter whether, in hallucinating, things phenomenally look a fully determinate way to the subject. It is not implausible that they do, nor is it implausible that the subject's neural state determines that they do. So suppose that the subject's neural state does determine a fully determinate way things look to the hallucinating subject. In that case, it seems reasonable that the subject's neural state also determines the fully determinate shade of pink that the hallucinated particular is, in the hallucination.

Regarding the second objection: If I see an Orchid Pink color chip only in the periphery of my visual field, I will be unable to tell what shade of pink it is. This need not be because I do not see the shade of the chip. The reason may be that my recognitional capacities are poor when I only have peripheral vision to go on. The same can be said of hallucinations. If a subject hallucinates something's being, in the hallucination, Orchid pink and in the periphery of the visual field, the subject will be unable to tell what shade of pink it is. This

need not be because the subject does not hallucinate the thing's being, in the hallucination, some determinate shade. The reason may be that the subject's recognitional capacities are poor when the subject only has peripheral visual hallucination to go on.

Regarding the third objection: It seems perfectly correct to suppose that the phenomenal character of an experience is influenced both by what the subject is experientially aware of and by the way things seem to the subject to be. Consider, for example, perceptually experiencing a Necker cube drawing. In experiencing the drawing, one is perceptually aware of a single fixed arrangement of lines. Nevertheless, the experience can have two distinct phenomenal characters. If we suppose that the way things seem to the subject to be can influence the phenomenal character of experience, this duality can be explained: One phenomenal character is due to an experience had when the lower left face of the cube seems to the subject to be in front. The other is due to an experience had when the upper right face of the cube seems to the subject to be in front.

Doubt 6. Suppose A has a visual hallucination of some kind K, and that in virtue of this fact A counts as visually hallucinating some merely hallucinated particular o. Plausibly, that A has a hallucination of kind K locally and metaphysically supervenes on the fact that A is in a neural state of some kind, N. But a subject A can also get into a state of kind N as a standard causal consequence of having A*'s retinas stimulated by light reflected off some real particular o*. In such a case, A* would be perceiving o*. Still, since A* would be in a state of kind N, and since N is the supervenience base for visual hallucinations of kind K, A* would be hallucinating o. So it emerges A* would be experiencing both o* and o. But this is surely absurd. So hallucinations involve no experiential awareness of merely hallucinated particulars.*

Reply. The doubt relies on the premise that hallucinations supervene on the hallucinator's neural state. But this premise is false, since there can be neural duplicates, one of whom is perceiving and not hallucinating, while the other is hallucinating and not perceiving. So, whether A is hallucinating supervenes not just on A's neural state, but on other factors as well - such as the causes for which A's is in A's neural state.

Though this reply is enough to undercut the argument for doubt 6, it should not be taken to suggest that no interesting mental states supervene on one's neural state. In particular,

there are views [e.g., Block's (1990)] on which the way things look to *A* does supervene on *A*'s neural state. This, of course, presupposes that the way things look to *A* is insensitive to the causes for which *A* is in *A*'s neural state. But such presupposition is not untenable, since things can look the same way to *A*, whether *A* is hallucinating or perceiving.

5.0 ON LEARNING FROM HALLUCINATION AND ILLUSION

Chapter 4 argued that there are such things as merely hallucinated particulars and hallucination-dependent states of affairs (henceforth: events). It further argued (in section 4.3.4) that the semantics of “to hallucinate” are existentially transparent when the verb is complemented with a phrase denoting a non-thinkable. The present chapter will build upon these foundations to provide a general theory of learning from experience, both perceptual and otherwise.

To remind the reader, the core of SP (presented in chapter 3) consists of the following claims:

(SP1) A perceiver, A , perceives x just in case A sees x , or smells x , or hears x , or ...

(SP2a) If no atomic experiential demonstrative concept occurs in a proposition of the form \mathcal{P} , then $C(\mathcal{P}, K)$ is a capacity to tell iff

- (i) the capacity can be successfully exercised,
- (ii) any successful exercise of the capacity by a subject A culminates in A 's non-inferentially coming to know some proposition p of the form \mathcal{P} owing to some (on-line or remembered) experiences E of the kind K , and
- (iii) any exercise of the capacity by a subject A culminates in A 's non-inferentially coming to believe some proposition p determined by (a) its having the form \mathcal{P} and (b) the circumstances under which A performs the exercise.

(SP2b) If a single atomic experiential demonstrative concept occurs in a proposition of the form \mathcal{P} , then $C(\mathcal{P}, K)$ is a capacity to tell iff

- (i) the capacity can be successfully exercised,
- (ii) any successful exercise of the capacity by a subject A culminates in A 's non-inferentially coming to know some proposition p of the form \mathcal{P} owing to some (on-line or remembered) experiences E of the kind K , and
- (iii) any exercise of the capacity by a subject A culminates in A 's non-inferentially coming to believe some proposition p determined by (a) its having the form \mathcal{P} , (b) the circumstances under which A performs the exercise, and (c) the fact that an experiential tracking effort F controls the atomic experiential demonstrative concept D taking the demonstrative position in \mathcal{P} .

(SP3a) A non-inferentially comes to *know* that p owing to (on-line or remembered) experiences E iff for some capacity to tell $C(\mathcal{P}, K)$, A 's *successful* exercise of $C(\mathcal{P}, K)$ culminates in A 's coming to believe that p owing to E .

(SP3b) A non-inferentially comes to *justifiably believe* that p owing to (on-line or remembered) experiences E iff for some capacity to tell $C(\mathcal{P}, K)$, A 's *non-reckless* exercise of $C(\mathcal{P}, K)$ culminates in A 's coming to believe that p owing to E .

(SP3c) A non-inferentially comes to *rationally believe* that p owing to (on-line or remembered) experiences E iff A *non-recklessly* and without inference comes to believe that p owing to E .

The reader will have noticed that in (SP2a)-(SP3c), the term “experience”, rather than the terms “perception” or “perceptual experience”, is the one used. This is no accident. According to SP, one can have capacities to tell $C(\mathcal{P}, K)$ where the kind of experiences K includes not just perceivings (e.g., seeings and hearings) but also experiences of non-perceptual types (e.g., visual and auditory hallucinations). Furthermore, according to SP, one can engage in experiential tracking efforts that are guided (either entirely or only in part) by non-perceptual experiences. One might even succeed in experientially tracking a target that is not perceived but, e.g., hallucinated. So SP is already equipped with the general scope that would allow it to account for cases in which one (non-inferentially) judges knowledgeably, justifiably or rationally on the strength of experiences, both perceptual and non-perceptual. The only thing missing is a clarification of what (non-perceptual) experiences are. This chapter will provide just this clarification.

Section 5.1 will build on the work of chapter 4 in providing a relationalist characterization of experiences. It will be part of this characterization that illusions do not constitute experiences in the relevant sense of the term. This surprising consequence will be discussed in section 5.2. Section 5.3 will then conclude both this chapter and this entire essay.

5.1 ON HALLUCINATIONS AND OTHER EXPERIENCES

As was already noted, chapter 4 argued (i) that there are such things as merely hallucinated particulars and hallucination-dependent states of affairs, and (ii) that the semantics of “to hallucinate” are existentially transparent when the verb is complemented with a phrase

denoting a non-thinkable. This allows us to take up again the broadly Austinian spirit of section 3.1, and with it to suggest that we have a way of using the “*A* visually hallucinates *x*” construction that is analogous to our way of using the “*A* sees *x*” construction. This way of using the construction can be characterized thus:

1. On the intended use of the “*A* visually hallucinates *x*” construction, what one says in uttering “*A* visually hallucinates *x*” does not entail what can be expressed by uttering “*A* believes *x*”, “*A* is inclined to believe *x*”, “*A* understands *x*”, or “*A* knows *x*”. For example, “Bob visually hallucinates a kitten’s expanding to the size of a lion” does not entail what can be expressed by uttering “Bob believes (or, is inclined to believe, ...) a kitten’s expanding to the size of a lion”. Such an utterance is ill formed. But more importantly, “Bob visually hallucinates a kitten’s expanding to the size of a lion” does not even entail that Bob believes (or, is inclined to believe, ...) *that* a kitten is expanding to the size of a lion. Being sober minded, Bob might be under no doubt at all that no kitten has ever expanded to the size of a lion. This possibility is important to remember: One can hallucinate something while knowing full well that one is hallucinating it.
2. On the intended use of the “*A* visually hallucinates *x*” construction, what one says in uttering “*A* visually hallucinates *x*” cannot be true if *x* is something that can be thought. Compare: If I utter, “Bob visually hallucinates that a kitten is expanding to the size of a lion” what I thereby say can be true, despite the fact that one can also think (for some odd reason) that a kitten is expanding to the size of a lion. So this is not the way if using the “*A* visually hallucinates *x*” construction that I have in mind. What sorts of things *can* be the references of “*x*” when one uses the “*A* visually hallucinates *x*” construction in the intended way and in saying something true? A partial answer is this: “*x*” can refer to anything that *A* can see, including objects, properties, relations and events. And potentially “*x*” can refer to other non-thinkables as well.
3. Intended uses of the “*A* visually hallucinates *x*” construction are substitutionally transparent. This means that necessarily, if (i) using the construction in the intended way one utters “*A* visually hallucinates *x*” at a context *c*, and (ii) *x* is identical to *y*, then were one to have used the construction in the the intended way in uttering “*A* visually hallucinates *y*” (at *c*), one would have said something with the same truth value. For example,

“Bob visually hallucinates Napoleon” entails “Bob visually hallucinates the Emperor of the French from 1804 to 1815”.

4. Intended uses of the “ A visually hallucinates x ” construction are existentially transparent. This means that what one says in uttering “ A visually hallucinates x ” (where one uses the construction in the intended way) entails that x exists or obtains. Here “exists” and “obtains” should be read tenselessly, i.e., as short of “exists, has existed or will exist” and “obtains, has obtained or will obtain” respectively. For example, “ A visually hallucinates Napoleon” entails that A visually hallucinates something that exists, has existed or will exist. In this case, it entails this for Napoleon - an ordinary mind-independent particular. Similarly, “ A visually hallucinates a pink elephant’s flying” entails that A visually hallucinates something that obtains, has obtained or will obtain. In this case, it entails this for some hallucination-dependent event, consisting in some abstract mind-dependent particular’s being, in A ’s hallucination, a flying pink elephant.¹

Having fixed our eyes on this (I believe familiar) way of using the “ A visually hallucinates x ” construction, we can introduce the notion of a visual hallucinating thus: A *visual hallucinating* is a mental state that can be reported on by the outlined way of using the “ A visually hallucinates x ” construction. In the rest of this essay, let us use the “ A visually hallucinates x ” construction only in the visual-hallucinating-reporting way. And if A visually hallucinates x , let us call x the “target” of the visual hallucinating, or say that the visual hallucinating is a visual hallucinating “of” x . The notions of an auditory hallucinating, an olfactory hallucinating, a gustatory hallucinating, a tactile hallucinating etc. can be introduced in the same way, along with the subsidiary notions of “a target” “of” a hallucinating of any one of these kinds. Having done all this, we can clarify the notion of a *hallucinating* quite simply. We say that

(SP4) A hallucinator, A , hallucinates x just in case A visually hallucinates x , or auditorily hallucinates x , or olfactorily hallucinates x , or ...

We now add to this that a *hallucinating* is any mental state that can be reported on by saying, of some A and x , that A hallucinates x .

¹For more on this, see sections 4.2.1 and 4.3.4.

With (SP1) and (SP4) clarifying the notions of hallucination and perception, we are well on our way to a full clarification of the notion of an experience. It only remains to address three last would-be types of experience: (i) So-called “inner” sense experiences (e.g., pain), (ii) experiences of entities such as after images and phosphenes, and (iii) illusions.

Though I largely leave inner sense experiences to be treated elsewhere, I do believe that they constitute a special kind of perceptual experience. I say a little more about this in section 3.2. Until I say more, I shall simply ask the reader to accept that inner sense experiences can be understood as perceptions and as presentational mental states. If this is accepted, it should also be accepted that SP can accommodate learning that occurs on the strength of such experiences.

As for experiences of after images and phosphenes: On the one hand, after images and phosphenes do not fit comfortably into the mold of merely hallucinated particulars, since they can be experienced by someone even when her perceptual system is functioning normally, giving no indication of a hallucinatory state. On the other hand, all the arguments regarding merely hallucinated particulars from chapter 4 can be recast to show that after images, phosphenes and their like are abstract and mind-dependent things, which exist contingently, in virtue of the fact that a person or persons created them in experiencing them. Given these two considerations, it seems best to simply say that when we experience after images or phosphenes, we are presented with certain existing mind-dependent entities. Whether we are presented with these entities via perceiving or hallucinating them is something we need not decide on. In fact, we might even say that we bear a third kind of experiential relation to them - experiencing*, say. What is important is this: True reports of the form “A visually (or, auditorily, olfactorily, ...) experiences x ”, where x is an entity of kind at hand, are reports of presentational mental states. And accordingly, such reports—like reports on perceivings and hallucinatings—have the four usual features: (i) They do not entail what can be expressed by uttering “A believes x ”, “A is inclined to believe x ”, “A understands x ”, or “A knows x ”, (ii) “ x ” refers to a non-thinkable. (iii) they are substitutionally transparent, and (iv) they are existentially transparent. They thus raise no special issues for SP’s account

of learning from experience.²

This brings me to the topic of illusions. Section 5.2 argues that illusions are not presentational mental states, and therefore not experiences. Since illusions are not experiences, a theory of learning from experiences (e.g., SP) need not address the putative epistemic powers of illusions.

The fact that SP does not address the putative epistemic powers of illusions might seem to be a drawback of the theory. To dispel this impression, section 5.2 argues that any case of learning from an illusion will also be a case of learning from certain proper experiences, to which the illusion is only a response. So, even though SP does not address the putative epistemic powers of illusions, it does account for all the cases of learning from an illusion. The explanatory power of SP, therefore, does not fall from that of any rival theory on which illusions are experiences. Furthermore, since SP can account for all the cases without postulating that illusions have epistemic powers, one might reasonably question whether illusions so much as have epistemic powers at all.

If illusions are not experiences, the upshot of the last few paragraphs quickly becomes visible. The upshot is that the notion of an experience can be clarified thus:

(SP5) An experience is either a perceiving, a hallucinating, or (if we admit the further category) an experiencing*.³

With (SP5) in hand, the presentation of SP is concluded. The reader can now see that, as promised, SP provides an entirely general account of learning from experiences.

Before moving on, I should briefly illustrate how SP accounts for learning from non-perceptual experiences. To do so, I would like to reconsider one of the cases of learning from hallucination discussed in section 4.3.3. In this case, Macbeth knows that he is hallucinating when he hallucinates the merely hallucinated dagger, Dag. Macbeth then forms the hallucination based judgment, of Dag, that it is, in his hallucination, a dagger. Macbeth expresses this judgment by telling himself:

(9) That is, in my hallucination, a dagger.

²For more on afterimages, see Ian Phillips' work (2012), which persuasively argues on empirical grounds that afterimage experiences are subjectively indistinguishable from perceptions of projected light phenomena.

³Note that (SP5) is consistent with there being experiences which lack phenomenal characters, though it does not imply it.

Intuitively, Macbeth's judgment is a non-inferential, knowledgeable and hallucination based one. SP can accommodate this intuition thus: Macbeth is visually hallucinating a hallucination dependent event, viz., Dag's being, in Macbeth's hallucination, a dagger. Macbeth also has some capacity to tell $C(\mathcal{T}\mathcal{H}\mathcal{A}\mathcal{T} - \mathcal{I}\mathcal{S} - \mathcal{I}\mathcal{N} - \mathcal{M}\mathcal{Y} - \mathcal{H}\mathcal{A}\mathcal{L}\mathcal{L}\mathcal{U}\mathcal{C}\mathcal{I}\mathcal{N}\mathcal{A}\mathcal{T}\mathcal{I}\mathcal{O}\mathcal{N} - \mathcal{A} - \mathcal{D}\mathcal{A}\mathcal{G}\mathcal{G}\mathcal{E}\mathcal{R}$, *visual hallucinations of something's being, in the subject's hallucination, a dagger*). Plausibly, Macbeth has this capacity partly in virtue of having a related capacity to tell that something is a dagger upon seeing its being a dagger. When Macbeth makes the judgment that (9) expresses, his judgment is the culmination of his successful exercise of the former capacity. This judgment ($\mathcal{T}\mathcal{H}\mathcal{A}\mathcal{T} - \mathcal{I}\mathcal{S} - \mathcal{I}\mathcal{N} - \mathcal{M}\mathcal{Y} - \mathcal{H}\mathcal{A}\mathcal{L}\mathcal{L}\mathcal{U}\mathcal{C}\mathcal{I}\mathcal{N}\mathcal{A}\mathcal{T}\mathcal{I}\mathcal{O}\mathcal{N} - \mathcal{A} - \mathcal{D}\mathcal{A}\mathcal{G}\mathcal{G}\mathcal{E}\mathcal{R}$) is consequently a non-inferential, knowledgeable and hallucination based one.

It should be noted that certain conditions have to be satisfied for the Macbeth's exercise of the capacity to be successful. First, the hallucination owing to which Macbeth makes the judgment has to be of a kind appropriate for the exercise of the capacity. This condition is satisfied since the hallucination owing to which Macbeth judges is a visual hallucinating of something's being, in the subject's hallucination, a dagger. Second, Macbeth must not be reckless in his exercise of the capacity. This condition is satisfied since Macbeth had no defeaters for the judgment, and in fact knows that he is hallucinating. Third, Macbeth's exercise of the capacity must occur under epistemically favorable circumstances. This condition is satisfied partly because Macbeth does not mishandle the exercise, which draws on a hallucination of the appropriate kind. Finally, fourth, the experiential demonstrative concept *THAT* occurring in Macbeth's judgment must be controlled by some successful experiential tracking effort through which Macbeth tracked what, in his hallucination, is a dagger. This condition is satisfied because Macbeth indeed attended to, and experientially tracked, Dag.

Let us now return to the topic of illusions.

5.2 ON ILLUSIONS

SP is based on the idea that experiences are presentational mental states, i.e., mental states that can be reported on by constructions that have the four just rehearsed features. This

idea might be taken to suggest the following line of thought: “We all know that there are three (or four) broad kinds of experiences - perceivings (including “inner sense” perceivings), hallucinatings, illusions and perhaps experiencings*. So far, I have been given presentational accounts of only two (or three) of those. Illusions have been left out. So I am owed a presentational account of illusions as well.”

Although tempting, this line of thought is to be resisted. According to SP, illusions are *not* presentational mental states in the relevant sense, and therefore do not count as experiences. There is no need to fight over words here: In *some* sense of “experience”, it is of course true that illusions are paradigmatic experiences. The point is only that in SP’s regimented sense of “experience”, they are not. But this will only emerge once we have a clear grasp of what illusions are.

Sadly, grasping what illusions are is often hindered by the fact that philosophers of perception very rarely attempt to characterize illusions, and when they do attempt this, the characterizations are often inadequate (and often knowingly so). In the next few paragraphs I will examine the flaws in the standard philosophical characterization of illusions, and try to revise it until the true phenomenon is made clear.

The standard characterization [see, e.g., [Brewer \(2011\)](#) and [Crane \(2011\)](#)] is this: An illusion is a mental state consisting in a perceived target’s feeling (e.g., looking, sounding, smelling, ...) to one a way that it is not. The most serious problem with this characterization was pointed out by Travis ([2013b](#)). Here is an illustration of his point: If you eyeball a white wall illuminated by a red light source, it will look to you as white walls illuminated by red light sources do look to you. So (assuming you are not spectrum inverted) it will look white and illuminated by a red light source to you. Similarly, if you eyeball a red wall illuminated by a white light source, it will look to you as red walls illuminated by white light sources do look to you. So (again assuming no inversion) it will look red and illuminated by a white light source to you. The issue is, however, that a white wall illuminated by a red light source and a red wall illuminated by white light source are ringers for each other.⁴ This means that the two walls look *the same way* to you. Consequently, we can say something true

⁴Here I harmlessly abstract away from the existence of metamers. The abstraction can be avoided by an easy relativization of the notion of a ringer to a class of perceivers.

about either one of the two walls by uttering (i) “to you it looks white and illuminated by a red light source”, and by uttering (ii) “to you it looks red and illuminated by a white light source”. This is so because the two complements of the “looks” verb pick out a single way a wall might look to you. And now we have a serious problem. While the two complements pick out a single way a wall might *look* to you, they pick out very different ways it might *be*. (After all, a single wall cannot be both white and red, no matter how it is illuminated.) In other words, a single way something looks to one can be specified in various ways, each specification corresponding to its being a different way. Hence, there is no fact of the matter about whether a wall is the way that it looks to one.⁵ There is only a fact of the matter about whether a wall is some way that it has been *specified* to look.

Since there is no fact of the matter about whether a wall is the way that it looks to one, there is also no fact of the matter about whether it looks to one the way it is. So illusions cannot be characterized as mental states consisting in a perceived target’s feeling to one a way that it is not. This standard characterization must be revised.

The revision required is not hard to find. We simply need to say that an illusion is a mental state consisting in a perceived target’s (visually, audibly, olfactorily, ...) seeming to one *to be* a way that it is not. This characterization does not have the problem of its predecessor. It has two others.

The first problem is that a perceived target might seem to one to be a way that it truly is, and still its seeming to be that way would constitute an illusion. (These illusions are called “veridical illusions”.) One demonstration of this possibility is given in a case due to Johnston (2006): Suppose one of two twins, who are in fact the same height, seems to be a different height than the other. This is because she is dressed in a shirt with horizontal stripes. Also suppose the twins are walking around in an Ames room. The room is designed to make any object seem to be different heights in different locations in the room. As the twins take different positions in the room, they seem to be different relative heights. Still, there are pairs of positions in the room such that when the twins are viewed standing in those positions, they seem to be the same height. When they are in those positions, the

⁵Note well: There is no fact of the matter about this, even though there are facts of the matter about how the wall looks, and about how the wall is.

illusory effects of the horizontally striped shirt exactly offset those of the Ames room. Thus, when the twins are viewed standing in those positions, they seem to be a way that they are. Nonetheless, since it is a mere accident that they seem to be the same height, their seeming to be the same height intuitively constitutes an illusion.

The solution to the first problem is to revise the characterization of an illusion to this: An illusion is a mental state consisting in a perceived target's either (visually, audibly, olfactorily, ...) seeming to one to be a way that it is not, or *accidentally* (visually, audibly, olfactorily, ...) seeming to one to be a way that it is. But this revision still leaves us with the second problem.

The second problem is that one might undergo an illusion even when one is not perceiving anything at all.

To illustrate the second problem, suppose I hallucinate a (merely hallucinated) Müller-Lyer stimulus. What I hallucinate is therefore a (merely hallucinated) piece of paper, with (merely hallucinated) black lines of *equal* length on it, where one line ends with (merely hallucinated) arrow heads and the other ends with (merely hallucinated) arrow tails. At the same time, I perceive nothing at all.

On the one hand, since I hallucinate (merely hallucinated) lines of equal length, I also hallucinate the (hallucination dependent) event of the the lines' being, in my hallucination, of equal length. On the other hand, since I hallucinate the lines' having, in my hallucination, the aforementioned elements on their ends, the lines will seem to me to be different lengths. So the lines will seem to me to be a way that they are not, and furthermore, a way that, in my hallucination, they are not. Intuitively, this means that I am undergoing an illusion. And yet, I am undergoing this illusion as I perceive nothing. So *pace* dominant opinions, illusions do not require object perception.

This second problem suggests a simplification to our last characterization of an illusion: An illusion is a mental state consisting either in its (visually, audibly, olfactorily, ...) seeming to one that things are a way that they are not, or in its accidentally (visually, audibly, olfactorily, ...) seeming to one that things are a way that they are. This, I believe, is the way we should think of illusions.

To see the appeal of the last characterization of illusions, consider the idea that one of

the functions our cognitive systems have is this: To take our experiences (presentationally conceived) as input, and to output non-accidentally true seemings—where x is a seeming iff for some p , x is the mental state that consists in its (visually, audibly, olfactorily, ...) seeming to one that p . When our cognitive systems are operating under favorable circumstances, they produce successful seemings. In a successful seeming, things non-accidentally seem to us to be ways that they are. But under unfavorable circumstances, our systems produce defective seemings. An illusion is simply a defective seeming. Sometimes a seeming is defective and still, in it, things seem to be a way that they are. In such cases one has a veridical illusion. At other times, a seeming is defective and in it things seem to be a way that they are not. In such cases one has a falsidical illusion.

To recapitulate: Illusions are defective seemings. A mental state x is a seeming iff for some p , x consists in its (visually, audibly, olfactorily, ...) seeming to one that p .

It immediately follows from the aforementioned that seemings (and therefore, illusions) are representational and not presentational mental states. They are representational states because they are evaluable for truth (veridicality, accuracy) at possible worlds. That p is something that can be (or, can fail to be) the case at a possible world w . Hence, when it seems to A that p , things either seem to A to be a way they are at w , or a way they are not at w . In the first case, the seeming is veridical at w . In the second case, it is falsidical at w .

Seemings are not presentational states (in SP's sense) because they lack three of the four characteristics presentational states have. First, if it seems to A that p , the phrase "that p " designates a thinkable thing. For example, if it seems to me that the Müller-Lyer lines are different lengths, I can also believe that they are different lengths. Second, reports of the form "it seems to A that p " are not substitutionally transparent. For example, it can seem to me that over there is Walt Disney without its seeming to me that over there is the man who won the most Oscars to date. Finally, reports of the form "it seems to A that p " are not existentially transparent. For a counterexample, reconsider a case from section 3.3:

Suppose I see two bricks - a red one and a green one - flying towards each other. As they hit they exchange colors - the red one turns green and the green one turns red. I can't hear the impact because I have wax in my ears. Immediately after the bricks hit, they oddly start flying back away from each other. I see this too.

Supposing all this happens against the appropriate background, it seems to me that a con-

stantly colored red brick is flying past a constantly colored green brick. But there exists nothing which seems to me to be a constantly colored red brick flying past a constantly colored green brick. In particular, neither brick seems to me to be this way. The reason is that before the two bricks hit, one of them seems to be about to fly past the other, while after they hit the other seems to have flown past the first.

It turns out, then, that seemings (and therefore, illusions) are not presentational mental states. From this it follows that they are also not experiences in SP's sense. There is therefore no reason to deny (SP5). Furthermore, since SP only claims to be a theory of knowledgeable (justified, or rational) *experience*-based belief formation, there is no reason for SP to address the putative epistemic powers of seemings.

I expect that some would now object, and say that SP's silence with respect to seemings is somewhat problematic. But such an objections should be resisted. For on the assumption that we can knowledgeably (justifiably, rationally) form beliefs on the strength of seemings, SP can account for all the cases in which we do so.

Suppose McX wants to say that one *can* non-inferentially acquire experience-based knowledge (justified beliefs, rational beliefs) owing to seemings. McX still has to accept, however, that seemings are the outputs of systems whose inputs are experiences (now construed strictly presentationally). So by McX's own lights, we non-inferentially acquire experience-based knowledge (justified beliefs, rational beliefs) in the following 2-step way: First, it seems to *A* that *p* because *A* has experiences *E* and because certain cognitive systems of *A*'s respond to *E* in a certain way. Second, *A* non-inferentially comes to know (justifiably believe, rationally believe) that *q* owing to its seeming to *A* that *p*. But a proponent of SP can now piggyback on McX's view. This proponent would simply stipulate that *A*'s capacities to tell are such that the entire process—starting with the cognitive processing of *E* and ending with *A*'s (knowledgeable, justified or rational) belief that *q*—counts as an exercise of an appropriate one of them; and that the resulting belief that *q* is therefore had owing to *E*. Thus, any case McX would like to classify as one of seeming-based (knowledgeable, justified or rational) belief formation, would also count as a case of experience-based (knowledgeable, justified or rational) belief formation that SP can accommodate.

To make the last argument more concrete, consider three examples of how SP would

accommodate cases that one might take to be ones of illusion-based belief formation:

(1) When I look at a white wall illuminated by a red light source, though I do see the wall, I do not see the wall's being white. That event is occluded by the red light striking the wall. What I do see is the white wall's reflecting of red light. And my seeing the white wall's reflecting of red light leads to its seeming to be red to me. Of course, the wall's being illuminated by a red light source constitutes unfavorable circumstances for the operation of my cognitive system. It is therefore no surprise that I end up having a defective seeming and being under a falsidical illusion. Still, if I end up justifiably believing that the wall is red, SP can account for this. By SP's lights, my belief is had owing to a non-reckless but unsuccessful exercise of a capacity to tell red walls upon seeing them. The exercise of this capacity is unsuccessful because it occurs under epistemically unfavorable circumstances (the wall's being illuminated by a red light source), which manifest the the fallibility of the capacity.




Figure 1: Kanizsa triangle

(2) When I look at the Kanizsa triangle stimulus, I see three pacman-shaped black circles and three black wedge-shaped lines. I also see these objects' being arranged in a certain way on a white piece of paper. Finally, I see a certain white triangular area at the center of the piece of paper. These seeings lead to that triangular area's seeming to me to be lying on top of three black circles. Of course, the arrangement of the various objects on the paper constitutes unfavorable circumstances for the operation of my cognitive system. My seeming is therefore defective and constitutes a falsidical illusion. Still, if I end up justifiably believing that something is on top of three black circles, SP can account for this. By SP's

lights, my belief is had owing to a non-reckless but unsuccessful exercise of a capacity to tell that something is on top of three black circles upon seeing its being so. The exercise of this capacity is unsuccessful because it occurs under epistemically unfavorable circumstances (items' being unusually arranged on a piece of paper), which manifest the the fallibility of the capacity.

(3) Suppose I look at a waterfall for a while, and then at a stationary object. I will then see the object's being stationary. I will also experience* (or see, or hallucinate) an after-image of motion's being superimposed over the object. These experiences lead to the stationary object's seeming to be both moving and stationary to me. Of course, the after-image's being superimposed over the object constitutes unfavorable circumstances for the operation of my cognitive system. My seeming is therefore defective and constitutes a (necessarily) falsidical illusion. Unlike the previous two cases, however, in this case I cannot even rationally believe, of the object, that it is both moving and stationary. SP can account for this as well. First, one can have no capacity to tell whose exercise can only relate one to necessarily false propositions, since such capacities cannot be successfully exercised. Second, even though one can have a capacity to form mere beliefs whose exercise relates one to necessarily false propositions, if the propositions are also manifestly false (as are propositions of the form *THAT I - I - BOTH - MOVING - AND - STATIONARY*), then all exercises of the capacity would be reckless.

I will end this section by making two final points:

First, I have just argued that, on the assumption that we can knowledgeably (justifiably, rationally) form beliefs on the strength of seemings, SP can account for all the cases in which we do so. But this is not to say that I accept that we can knowledgeably (justifiably, rationally) form beliefs on the strength of seemings. In fact, I seriously doubt this.

My doubts are based on two grounds. The first is that the best account I know of learning from seemings is the version of epistemic representationalism (ER) we encountered in section 2.5. As I argued there, I believe this account fails in certain cases. My second ground for doubt has to do with considerations of simplicity and intuitiveness. Any account on which we knowledgeably (justifiably, rationally) form beliefs on the basis of seemings will be a 2-step account, like McX's. This seems needlessly complicated. SP makes do with a

simpler one-step process. Furthermore, SP replaces McX’s theoretically laden transitions from seemings to beliefs, with transitions from experiences to beliefs that are—as section 3.1 showed—a part of our implicit, intuitive, worldview. In short, SP cuts the middleman (i.e., the seemings) out of the epistemological story. This intuitive simplification is reason enough to doubt whether seemings belong in the epistemological story in the first place.

Second, some readers might still wonder why I treat seemings and experiences as ontologically different. “You are willing to postulate certain entities (e.g., the merely hallucinated dagger, Dag) whose existence depends on the existence of certain experiences”, they will remind me. “But if you are willing to spend your ontological capital in this extravagant way, why not do the same with seemings? Why not postulate certain *seeming dependent* entities, to parallel your experience dependent entities?”

My response is that postulating experience dependent entities is not ontologically extravagant, but postulating seeming dependent entities would be. Experience dependent entities are not extravagant because, as chapter 4 argued, there are intuitive, phenomenological, epistemological and semantic reasons to postulate them. I know of no similar reasons to postulate seeming dependent entities.

To illustrate how seeming dependent entities would be idle, consider a familiar toy example: I look at a white wall illuminated by a red light source, so that the wall seems to be red to me. Should we postulate such a seeming dependent entity as *the seeming color of the wall*? Let’s see. (i) The seeming color of the wall would be superfluous in any list of entities that look or seem to be some way to me. The entities that look or seem to be some way to me are all enumerated in the list of entities that I experience. This list includes the following: the wall, its various parts, the color of the light it reflects, and its reflecting of red light. Once I have specified how all those entities look or seem to be to me, there seems to be little reason to additionally insist that there are ways that the seeming color of the wall looks or seems to be to me. Better accept then that the seeming color of the wall neither looks nor seems to be any way to me.

(ii) The seeming color of the wall would be superfluous in any list of entities that I can attend to, experientially track, refer to with an experiential demonstrative or form beliefs about. These “epistemically engageable” entities are all enumerated in the aforementioned

list of entities that I experience. In particular, if I attend, track and come to believe of some property that it is the color red, that property would plausibly be the color of the light that the wall reflects. Once this is recognized, there seems to be little reason to additionally insist that the seeming color of the wall is an epistemically engageable entity. Better accept then that the seeming color of the wall is not epistemically engageable.

(iii) Experienced entities are semantically useful. In particular, and as I noted in sections 3.2, 4.3.4 and 5.1, the truth of certain sentences of the forms “*A* experiences *x*”, “*A* hears *x*”, “*A* visually hallucinates *x*” etc. seems to require the existence of experienced entities, including experience dependent ones. But I know of no similar semantic use for seeming dependent entities.

(iv) In chapter 4 I argued that postulating hallucination dependent entities is part of our intuitive worldview. The argument (which can be generalized for all experience dependent entities) relied on showing that we believe certain things about hallucination dependent entities. For example, we believe that we create such entities, that they are different from ordinary entities, that they can terrify us, that we can worship them, that they can become famous, etc. But I doubt that we have similar beliefs about seeming dependent entities. So I doubt that seeming dependent entities are a part of our intuitive worldview.

It emerges that seeming dependent entities are not intuitive, semantically useful or epistemically engageable. It also seems that they neither look nor seem to be any way to us. This casts doubt over whether we should accept that they exist. Of course, there may be good reasons to accept their existence that I am simply not considering. But if so, let those reasons be given. Until that time, postulating seeming dependent entities would seem to be ontologically extravagant.

5.3 CONCLUSION

In the course of this essay we touched upon a wide range of topics - what perceptions, illusions and hallucinations are, how experiential demonstrative reference is effected, the possibility and ramifications of inverted spectrum scenarios, and the epistemology of experience-based

knowledge and belief. No doubt there is much more to be said on each of them. But this essay was never meant to exhaust any of these issues. Rather, it was meant to motivate a very abstract, but hopefully also a very fruitful, epistemological framework - SP. It tried to motivate it by showing that the framework has the following advantages:

1. SP is a mild rational reconstruction of our implicit conception of non-inferential perception based knowledge acquisition.
2. SP accounts not only for non-inferential perception based knowledge acquisition, but more generally for non-inferential experience based knowledge acquisition.
3. SP accounts not only for experiential knowledge, but also for experience based justified beliefs and for experience based rational beliefs.
4. SP accounts for learning on the strength of both on-line experiences, and remembered experiences.
5. SP does not require that truth (veridicality, accuracy) conditions be assigned to experiences, so it avoids the pitfalls of representational accounts of non-inferential experience based knowledge acquisition presented in chapter 2.
6. SP offers distinctively simple and compelling accounts of experiential tracking and experiential demonstrative concept use.
7. SP can accommodate the fact that it is possible for people whose color spectra are inverted (relative to each other) to both knowledgeably and directly judge the colors of things.

I hope that these advantages are enough to persuade the reader that SP is well worth further study. I also hope that they are enough to motivate a broadly relationalist conception of experience; and in particular, one that excludes illusions from the experiential domain.

6.0 BIBLIOGRAPHY

- Alston, W. P. (1999). Back to the theory of appearing. *Philosophical Perspectives*, 13:181–203.
- Austin, J. L. (1970). Other minds. In Urmson, J. O. and Warnock, G. J., editors, *Philosophical Papers*, chapter 4, pages 76–116. Clarendon Press, Oxford, 2 edition.
- Block, N. (1986). Advertisement for a semantics for psychology. *Midwest Studies in Philosophy*, 10:615–678.
- Block, N. (1987). Functional role and truth conditions. *Proceedings of the Aristotelian Society*, 61:157–181.
- Block, N. (1990). Inverted earth. *Philosophical Perspectives*, 4:53–79.
- Block, N. (2007a). Mental paint. In *Consciousness, Function and Representation: Collected papers*, volume 1, chapter 24, pages 533–570. MIT Press, Cambridge MA.
- Block, N. (2007b). Sexism, racism, ageism, and the nature of consciousness. In *Consciousness, Function, and Representation*, volume 1, chapter 25, pages 571–601. MIT Press, Cambridge, MA.
- Block, N. (2007c). Wittgenstein and qualia. *Philosophical Perspectives*, 21:73–115.
- Boynton, R. M. (1979). *Human Color Vision*. Holt Rinehart and Winston, New York.
- Brewer, B. (2011). *Perception and its Objects*. Oxford University Press, Oxford.

- Broackes, J. (2007). Black and white and the inverted spectrum. *The Philosophical Quarterly*, 57(227):161–175.
- Brock, S. (2002). Fictionalism about fictional characters. *Noûs*, 36:1–21.
- Burge, T. (1979). Individualism and the mental. *Midwest Studies in Philosophy*, 4:73–121.
- Burge, T. (1986). Individualism and psychology. *The Philosophical Review*, 95(1):2–45.
- Burge, T. (2003a). Perceptual entitlement. *Philosophy and Phenomenological Research*, 67:503–548.
- Burge, T. (2003b). Qualia and intentional content: Reply to Block. In Hahn, M. and Ramberg, B., editors, *Reflections and Replies - Essays on the Philosophy of Tyler Burge*, chapter 21, pages 405–415. MIT Press.
- Burge, T. (2010). *Origins Of Objectivity*. Oxford University Press.
- Byrne, A. and Hilbert, D. R. (1997). Colors and reflectances. In Byrne, A. and Hilbert, D. R., editors, *Readings on Color*, volume 1, chapter 14, pages 263–288. MIT Press.
- Byrne, A. and Hilbert, D. R. (2003). Color realism and color science. *Behavioral and Brain Sciences*, 26:3–64.
- Campbell, J. (2002). *Reference and Consciousness*. Oxford University Press, Oxford.
- Campbell, J. (2009). Consciousness and reference. In McLaughlin, B., Beckermann, A., and Walter, S., editors, *Oxford Handbook of Philosophy of Mind*, pages 648–662. Oxford University Press, Oxford.
- Campbell, J. and Cassam, Q. (2014). *Berkeley's Puzzle*. Oxford University Press, Oxford.
- Cavell, S. (1979). *The Claim of Reason*. Oxford University Press, Oxford.
- Chalmers, D. J. (2010a). Perception and the fall from Eden. In *The Character of Consciousness*, pages 381–454. Oxford University Press.

- Chalmers, D. J. (2010b). The representational character of experience. In *The Character of Consciousness*, pages 339–380. Oxford University Press.
- Conduct, M. (2012). Naive realism without disjunctivism about experience. *Consciousness and Cognition*, 21(2):727–736.
- Crane, T. (2011). The problem of perception. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Spring 2011 edition.
- Crane, T. (2013). *The Objects of Thought*. Oxford University Press, Oxford.
- Currie, G. (1990). *The Nature of Fiction*. Cambridge University Press, Cambridge.
- Davies, M. (1992). Perceptual content and local supervenience. *Proceedings of the Aristotelian Society*, 92:21–45.
- Davies, M. (1997). Externalism and experience. In Block, N., Flanagan, O., and Güzeldere, G., editors, *The Nature of Consciousness: Philosophical Debates*, chapter 16, pages 309–327. MIT Press, Cambridge, MA.
- Dennett, D. C. (1988). Quining qualia. In Marcel, A. J. and Bisiach, E. E., editors, *Consciousness in Contemporary Science*, pages 42–77. Oxford University Press, Oxford.
- Dennett, D. C. (1991). *Consciousness Explained*. Back Bay books, New York.
- Dretske, F. (1995). *Naturalizing the Mind*. MIT Press, Cambridge, MA.
- Evans, G. (1982). *The Varieties of Reference*. Oxford University Press, Oxford.
- Everett, A. (2005). Against fictional realism. *The Journal of Philosophy*, 102:624–649.
- Everett, A. (2013). *The Nonexistent*. Oxford University Press, Oxford.
- Fine, K. (1984). Critical review of Parsons' *Non-Existent Objects*. *Philosophical Studies*, 45:95–142.
- Fish, W. (2009). *Perception, Hallucination, and Illusion*. Oxford University Press, Oxford.

- Fodor, J. A. (1992). A theory of content. In *A Theory of Content and Other Essays*. MIT Press.
- Genone, J. (2014). Appearance and illusion. *Mind*, pages 1–38. First published online: June 5, 2014.
- Gupta, A. (2006). *Empiricism and Experience*. Oxford University Press.
- Gupta, A. (2012). An account of conscious experience. *Analytic Philosophy*, 53(1):1–29.
- Hardin, C. L. (1997). Reinventing the spectrum. In Byrne, A. and Hilbert, D. R., editors, *Readings on Color*, volume 1, chapter 15, pages 289–301. MIT Press.
- Harman, G. (1982). Conceptual role semantics. *Notre Dame Journal of Formal Logic*, 23:242–256.
- Harman, G. (1987). (Nonsolipsistic) conceptual role semantics. In LePore, E., editor, *New Directions in Semantics*. Academic Press, London.
- Harman, G. (1990). The intrinsic quality of experience. *Philosophical Perspectives*, 4:31–52.
- Harrison, B. (1973). *Form and Content*. Blackwell, Oxford.
- Hilbert, D. R. and Kalderon, M. E. (2000). Color and the inverted spectrum. In Davis, S., editor, *Color Perception: Philosophical, Psychological, Artistic, and Computational Perspectives*, chapter 12, pages 187–214. Oxford University Press.
- Hill, C. S. (2009). *Consciousness*. Cambridge University Press.
- Hinton, J. M. (1967). Visual experiences. *Mind*, 76(302):217–227.
- Hurvich, L. M. (1982). *Color Vision*. Sinauer Associates Inc., MA.
- Inwagen, P. V. (1977). Creatures of fiction. *American Philosophical Quarterly*, 14:299–308.
- Johnston, M. (2004). The obscure object of hallucination. *Philosophical Studies*, 120:113–183.

- Johnston, M. (2006). Better than mere knowledge? The function of sensory awareness. In Gendler, T. S. and Hawthorne, J., editors, *Perceptual Experience*, pages 260–290. Oxford University Press.
- Kay, P. (1999). Asymmetries in the distribution of composite and derived basic color categories. *Behavioral and Brain Sciences*, 22(6):957–958.
- Kriegel, U. (2002). Phenomenal content. *Erkenntnis*, 57(2):175–198.
- Kripke, S. A. (1980). *Naming and Necessity*. Harvard University Press.
- Kripke, S. A. (2013). *Reference and Existence*. Oxford University Press.
- Kroon, F. (1994). A problem about make-believe. *Philosophical Studies*, 75:201–229.
- Kroon, F. (2011). The fiction of creationism. In Lihoreau, F., editor, *Truth In Fiction*, volume 38 of *Philosophical Analysis*, pages 203–221. Ontos Verlag, Frankfurt.
- Lewis, D. (1978). Truth in fiction. *American Philosophical Quarterly*, 15:37–46.
- Lycan, W. G. (1996). *Consciousness and Experience*. MIT Press, Cambridge, MA.
- Martin, M. G. F. (2002). Particular thoughts and singular thought. In O’Hear, A., editor, *Logic, Thought and Language*, volume 51 of *Royal Institute of Philosophy Supplement*, pages 173–214. Cambridge University Press.
- Martin, M. G. F. (2004). The limits of self-awareness. *Philosophical Studies*, 120:37–89.
- McDowell, J. (1996). *Mind and World*. Harvard University Press, Cambridge, MA, 2 edition.
- McDowell, J. (1998). The content of perceptual experience. In *Mind, Value and Reality*, chapter 16, pages 341–358. Harvard University Press, Cambridge, MA.
- McDowell, J. (2009). Avoiding the myth of the given. In *Having the World in View*, chapter 14, pages 256–272. Harvard University Press, Cambridge, MA. Originally published 2008.

- McGinn, C. (1991). *The Problem of Consciousness*. Blackwell, Oxford.
- Meinong, A. (1960). The theory of objects. In Chisholm, R. M., editor, *Realism and the Background of Phenomenology*, chapter 4, pages 76–117. The Free Press, New York. Translated by Isaac Levi, D. B. Terrell, and Roderick M. Chisholm.
- Millar, A. (2008). Perceptual-recognitional abilities and perceptual knowledge. In Haddock, A. and Macpherson, F., editors, *Disjunctivism: Perception, Action, Knowledge*, pages 330–347. Oxford University Press.
- Millikan, R. G. (2004). *Varieties of Meaning*. MIT Press.
- Nida-Rümelin, M. (1996). Pseudonormal vision: An actual case of qualia inversion? *Philosophical Studies*, 82(2):145–157.
- Palmer, S. E. (1999). Color, consciousness, and the isomorphism constraint. *Behavioral and Brain Sciences*, 22(6):923–989.
- Parsons, T. (1980). *Nonexistent Objects*. Yale University Press, New Haven.
- Pautz, A. (2006). Sensory awareness is not a wide physical relation: An empirical argument against externalist intentionalism. *Nouûs*, 40(2):205–240.
- Pautz, A. (2007). Intentionalism and perceptual presence. *Philosophical Perspectives*, 21:495–541.
- Pautz, A. (2010). Why explain visual experience in terms of content? In Nanay, B., editor, *Perceiving the World*, chapter 10, pages 254–309. Oxford University Press, Oxford.
- Peacocke, C. (1992). *A Study of Concepts*. MIT Press.
- Peacocke, C. (2001). Does perception have a nonconceptual content? *The Journal of Philosophy*, 98(5):239–264.
- Peacocke, C. (2004). *The Realm of Reason*. Oxford University Press, Oxford.

- Phillips, I. (2012). Afterimages and sensation. *Philosophy and Phenomenological Research*, 87:417–453.
- Piantanida, T. P. (1974). A replacement model of x-linked recessive colour vision defects. *Annals of Human Genetics*, 37(4):393–404.
- Priest, G. (2005). *Towards Non-Being*. Oxford University Press, Oxford.
- Pryor, J. (2000). The skeptic and the dogmatist. *Noûs*, 34(4):517–549.
- Putnam, H. (1975). The meaning of 'meaning'. In *Mind, Language and Reality: Philosophical Papers*, volume 2. Cambridge University Press, Cambridge.
- Roessler, J. (2009). Perceptual experience and perceptual knowledge. *Mind*, 118(472):1013–1041.
- Sainsbury, R. M. (2010). *Fiction and Fictionalism*. New Problems of Philosophy. Routledge, New York.
- Salmon, N. (1998). Nonexistence. *Noûs*, 32:277–319.
- Schiffer, S. (1996). Language-created language-independent entities. *Philosophical Topics*, 24:149–167.
- Searle, J. (1983). *Intentionality: An Essay in the Philosophy of Mind*. Cambridge University Press.
- Searle, J. R. (1975). The logical status of fictional discourse. *New Literary History*, 6:319–332.
- Shoemaker, S. (1982). The inverted spectrum. *The Journal of Philosophy*, 79(7):357–381.
- Shoemaker, S. (1994). Phenomenal character. *Noûs*, 28(1):21–38.
- Shoemaker, S. (2007). A case for qualia. In McLaughlin, B. P. and Cohen, J., editors,

- Contemporary Debates in the Philosophy of Mind*, chapter 18, pages 319–332. Blackwell, Malden, MA.
- Siegel, S. (2010). *The Contents of Visual Experience*. Oxford University Press.
- Smith, A. D. (2002). *The Problem of Perception*. Harvard University Press, Cambridge, MA.
- Snowdon, P. (1980-1981). Perception, vision and causation. *Proceedings of the Aristotelian Society*, 81:175–192.
- Snowdon, P. (1990). The objects of perceptual experience. *Proceedings of the Aristotelian Society*, 64:121–150.
- Taylor, D. M. (1966). The incommunicability of content. *Mind*, 75(300):527–541.
- Thomasson, A. L. (1999). *Fiction and Metaphysics*. Cambridge University Press.
- Thomasson, A. L. (2010). Fiction, existence and indeterminacy. In Woods, J., editor, *Fictions and Models: New Essays*, pages 109–148. Philosophia Verlag.
- Travis, C. (2013a). *Perception: Essays After Frege*. Oxford University Press, Oxford.
- Travis, C. (2013b). The silences of the senses. In *Perception: Essays After Frege*, chapter 1, pages 23–58. Oxford University Press, Oxford.
- Tye, M. (1995). *Ten Problems of Consciousness*. MIT Press, Cambridge, MA.
- Tye, M. (2000). *Consciousness, Color and Content*. MIT Press, Cambridge, MA.
- Tye, M. (2009). *Consciousness Revisited: Materialism without Phenomenal Concepts*. MIT Press, Cambridge, MA.
- Walton, K. L. (1990). *Mimesis as Make-Believe*. Harvard University Press, Cambridge, MA.