

A DEFENSE OF LIBERALISM IN THE EPISTEMOLOGY OF PERCEPTION

By

MEGAN FEENEY

A dissertation submitted to the

School of Graduate Studies

Rutgers, The State University of New Jersey

In partial fulfillment of the requirements

For the degree of

Doctor of Philosophy

Graduate Program in Philosophy

Written under the direction of

Susanna Schellenberg

And approved by

---

---

---

---

---

New Brunswick, New Jersey

October 2019

## ABSTRACT OF THE DISSERTATION

### A Defense of Liberalism in the Epistemology of Perception

By MEGAN FEENEY

Dissertation Director:

Susanna Schellenberg

I endorse liberalism, the view that we can be immediately perceptually justified in believing at least some propositions about the external world. In the dissertation, I defend liberalism against several contemporary challenges.

In Chapter 1, I address the scope of immediate perceptual justification. I argue that we can be immediately perceptually justified in believing propositions that ascribe kind properties to perceived individuals, (e.g., *that's an avocado*), and that this justification is grounded in recognitional competence. By appeal to recognitional competence, we can defuse the threat posed by McGrath's (2017) Only-Because Argument, which if successful, would restrict the scope of immediate perceptual justification to beliefs about how objects look.

Many liberal views appeal to seemings, but there is little consensus about what seemings are. In Chapter 2, I develop what I call the epistemic feeling account of seemings. Epistemic feelings are a class of phenomenally conscious mental states that motivate epistemic performance; familiar examples include the tip of the tongue feeling, and the eureka experience or aha! moment. I highlight an epistemic feeling—called the feeling of rightness—that cognitive psychologists have used to explain when subjects are

likely to simply endorse the output of some Type-1 process in judgment without engaging in Type-2 reasoning. I characterize the content, phenomenal character, and functional role of feelings of rightness and argue that they are well-suited to play the role of seemings in several central cases.

In Chapter 3, I address a third contemporary challenge for liberalism. White (2006) argues that Bayesian confirmation theory is in tension with liberalism, and since the former is such a fruitful paradigm, we have reason to reject liberalism. I consider three ways of reconstructing White's argument and argue that none succeeds. The most compelling version of the argument depends on intuitions that, in certain cases, the rational response to perceptual experience is to become more confident in skeptical hypotheses. I argue that these intuitions are driven by the presence of higher-order evidence in the cases, and that once we recognize this, the intuitions pose no threat to liberalism.

Finally, in Chapter 4, I return to considering the scope of immediate perceptual justification. On a natural view, the content of perceptual experience fully determines the scope of immediate perceptual justification. I highlight a class of cases—which I call *mismatch cases*—that show this natural thought to be incorrect. In mismatch cases, subjects are immediately perceptually justified in believing *p*, even when *p* is not among the contents of their perceptual experiences. After arguing that the cases should be characterized in this way, I consider the implications for liberalism. I argue that phenomenal conservatism struggles to handle mismatch cases by requiring a tight semantic connection between the content of experience and the content of beliefs that are immediately perceptually justified. Finally, I highlight two families of liberal views that

can account for mismatch cases. Permissive evidentialist views handle the cases by adopting a weakened evidential support relation. Basic method views handle them by explaining immediate perceptual justification by appeal to noninferential competences, methods, or processes. Finally, I briefly offer reasons for thinking that basic method views represent the more promising solution.

## Acknowledgements

I am grateful to Susanna Schellenberg, for spending countless hours reading and discussing my work, and for providing constant support and mentorship, beginning when I was a prospective student. I feel deeply fortunate to have had such a generous mentor.

Thanks also to Frankie Egan, Matt McGrath, Casey O’Callaghan, and Ernie Sosa for your guidance and feedback. It has been a tremendous privilege to learn from you all. I’m especially grateful for the ways you’ve shown, in your response to me and others, that you don’t have to sacrifice kindness to do rigorous philosophy. Special thanks to Susanna and Ernie for running such helpful dissertation groups over the years.

For comments and discussion, I’m grateful to Austin Andrews, Robert Audi, Austin Baker, Hanoch Ben-Yami, D. Black, Ben Bronner, Alex Byrne, Laura Callahan, Marilie Coetsee, Santiago Echeverri, Carolina Flores, Danny Forman, Georgi Gardiner, Alessandro Giovannelli, Veronica Gomez, James Gullledge, E.J. Green, Ben Henke, Chris Hill, Liz Jackson, Anton Johnson, Nico Kirk-Giannini, Nic Koziolk, Maria Lasonen-Aarnio, Ting-An Lin, Barry Loewer, Meghan Masto, Aidan McGlynn, Brian McLaughlin, Owen McLeod, Neil Mehta, Daniel Muñoz, Ram Neta, Olivia Odoffin, George Panichas, Dee Payton, Emily Prychitko, Howard Robinson, Pamela Robinson, Ezra Rubenstein, Joe Shieber, Eli Shupe, David Thorstad, Fritz Warfield, and Chris Willard-Kyle. Apologies to those I am forgetting.

Thanks also to audiences and commentators at the World Congress of Philosophy 2019, the Eastern APA 2018, the Rutgers-CEU Mind Conference 2018, Lafayette College, the UT Austin Graduate Conference 2017, the Notre Dame-Northwestern

Graduate Epistemology Conference 2017, and the University of Edinburgh Graduate Epistemology Conference 2017.

I could not have finished this project without the constant love, support, and good advice of Will Fleisher, Janet Ives, Nancy Ives, and Ed Wheeler. I dedicate the project in memory of my dad, Bill Feeney, who drove over 2,000 miles (roughly the distance across the continental US) to haul me to university visits as a kid.

Finally, some of the ideas in Chapter 3—Giving Liberals a Hand: A Defense of Liberalism Against Three Bayesian Arguments—served as the foundation of the following paper:

Feeney, M. & Schellenberg, S. (forthcoming). Bayesian liberalism. In K. M. Vogt & J. Vlasits (Eds.), *Epistemology after Sextus Empiricus*. Oxford: Oxford University Press.

# Table of Contents

<b>ABSTRACT</b> .....	ii
<b>ACKNOWLEDGEMENTS</b> .....	v
<b>1. INTRODUCTION</b> .....	1
<b>2. RECOGNITIONAL COMPETENCE AND KNOWING WHAT THINGS LOOK LIKE</b> .....	9
<i>THE ONLY-BECAUSE ARGUMENT</i> .....	10
<i>RECOGNITIONAL COMPETENCE</i> .....	18
<i>EXPLAINING COUNTERFACTUAL DEPENDENCE</i> .....	26
<i>INTERVENTIONIST COUNTERFACTUALS AND METAPHYSICAL DEPENDENCE</i> .....	32
<i>INTERVENING ON KNOWLEDGE VIA BELIEF</i> .....	37
<i>INTERVENING ON KNOWLEDGE VIA JUSTIFICATION</i> .....	42
<i>CONCLUSIONS</i> .....	50
<b>3. SEEMINGS ARE EPISTEMIC FEELINGS</b> .....	52
<i>TWO ACCOUNTS OF SEEMINGS</i> .....	53
<i>THE ONTOLOGICAL AND SEMANTIC THESES</i> .....	55
<i>WHAT ARE SEEMINGS?</i> .....	59
<i>EPISTEMIC FEELINGS</i> .....	68
<i>FEELINGS OF RIGHTNESS</i> .....	73
<i>CONCLUSIONS</i> .....	79
<b>4. GIVING LIBERALS A HAND: A DEFENSE OF LIBERALISM AGAINST THREE BAYESIAN ARGUMENTS</b> .....	81
<i>WHAT IS LIBERALISM?</i> .....	82
<i>THE BAYESIAN ARGUMENT AGAINST LIBERALISM</i> .....	89
<i>REJECTING THE CONTENT AND BELIEF ASSUMPTIONS</i> .....	93
<i>REVISING THE BAYESIAN ARGUMENT</i> .....	95
<i>HIGHER-ORDER EVIDENCE AND THE BAYESIAN ARGUMENT FROM CASES</i> .....	101
<i>CONCLUSIONS</i> .....	112

<b>5. THE SCOPE OF IMMEDIATE PERCEPTUAL JUSTIFICATION.....</b>	<b>114</b>
<i>MISMATCH CASES</i> .....	<i>116</i>
<i>THE MISMATCH ARGUMENT</i> .....	<i>125</i>
<i>TROUBLE FOR PHENOMENAL CONSERVATISM</i> .....	<i>136</i>
<i>TWO STRATEGIES FOR HANDLING MISMATCH</i> .....	<i>145</i>
<i>CONCLUSIONS</i> .....	<i>150</i>
<b>REFERENCES.....</b>	<b>152</b>

## Introduction

Liberalism, or the view that we can be immediately perceptually justified in believing propositions about the external world, is a popular and attractive alternative to classical foundationalism. According to the latter, we are only ever immediately justified in believing propositions about our own mental states, from which we must infer propositions about the external world. Liberalism is attractive because, at least at first glance, it lays the groundwork for a stronger anti-skepticism than classical foundationalism offers. The liberal is unburdened by the need to recover the full range of our knowledge of the external world from the skeptic by reconstructing inferences from the thin foundation of our introspective knowledge.

However, liberalism faces a number of serious contemporary challenges. Part of the anti-skeptical promise of liberalism depends on its potential to allow a wide range of propositions about the external world to count as immediately justified. Thus, the liberal faces the challenge of delimiting the scope of immediate perceptual justification, or of specifying in a principled way which types of propositions about the external world can be immediately perceptually justified. Liberals also face the challenge of integrating an account of immediate perceptual justification with Bayesian confirmation theory, in light of arguments—most prominently in White (2006)—that the two are incompatible with each other or at least in significant tension.<sup>1</sup> Finally, given the rise of liberal views that carve out a central role for *seemings*, many liberals face the challenge of characterizing this mental state and of clarifying the justification for recognizing seemings at all.

---

<sup>1</sup> White (2006) actually targets what I call Mooreanism, following Silins (2007). Mooreanism is a response to external world skepticism that depends on liberalism. I reconstruct White's argument so that it targets liberalism more directly.

In this project, I defend liberalism by developing an integrated response to each of these challenges. The dissertation is intended to be a collection of standalone papers. However, the arguments in each chapter reinforce the others.

First, a quick note on terminology. I avoid using the term “dogmatism” and instead, following Silins (2007), distinguish among three views that are commonly associated with dogmatism. Liberalism is the view that we can be immediately perceptually justified in believing at least some propositions about the external world. (Chapters 1, 3, and 4 each spend some time characterizing the distinction between immediate and mediate justification, but Chapter 4, Section 1 has the most extended discussion.) This is the view that I endorse. Phenomenal conservatism is a version of liberalism; one way to render phenomenal conservatism is as the principle that S has immediate *prima facie* justification to believe that p iff it seems to her that p. However, one can be a liberal without being a phenomenal conservative, as liberal views needn’t appeal to seemings at all, (e.g., Lyons 2009). (For discussion, see Chapter 3, Section 4). Finally, Mooreanism is the view that we can acquire justification for the first time to disbelieve skeptical hypotheses by performing Moorean inferences like the following: 1) *I have a hand*. 2) *If I have a hand, then I’m not a BIV*. 3) *I’m not a BIV*. Mooreans hold that we can be immediately perceptually justified in believing Premise 1. However, it is possible to be a liberal without being a Moorean. (For example, one can deny the closure principle.) I’m concerned here with defending liberalism, as distinct from phenomenal conservatism and Mooreanism.

In Chapter 1, I begin to address the scope challenge. I detail a notion of recognitional competence and use it to argue that we can be immediately perceptually

justified in believing propositions that ascribe kind properties to perceived individuals, (e.g., *that's an avocado*). Very roughly, a recognitional competence is a reliable disposition to believe truly that *o* is *F* (in the absence of defeaters and under relevant conditions) in response to perceiving that *o* has a certain collection of low-level properties.

This notion of recognitional competence helps us respond to a serious contemporary challenge to liberalism—McGrath's (2017) Only-Because Argument. McGrath points out that our perceptual knowledge that a perceived individual is an *F* depends on our knowledge of what *F*s look like. When I see an unlabeled avocado in the grocery store, for example, I wouldn't know that it's an avocado unless I knew what avocados looked like. This suggests that my perceptual knowledge that it's an avocado is mediated by my knowledge of what avocados look like, which, McGrath argues, is propositional knowledge. If the Only-Because Argument succeeds, it would show that perceptual knowledge of a perceived individual's kind properties must be mediated knowledge. Thus, the argument imposes a significant restriction on the scope of immediate perceptual justification.

The recognitional competence account offers the following solution. Sophisticated human subjects often possess both a competence to recognize *F*s and propositional knowledge that *F*s look *W*. However, as I argue, possessing a recognitional competence is neither necessary nor sufficient for possessing that piece of propositional knowledge. The truth of counterfactuals like *If I didn't know what avocados look like, I wouldn't know that this is an avocado* is best explained, not by holding that the latter knowledge must be based on the former. Rather, the key is that my knowledge of what

avocados look like, and my knowledge that this is an avocado, when grounded in recognitional competence, depend on each other in more subtle ways. For example, in an interventionist framework (Woodward 2003; Woodward & Hitchcock 2003), if we were to manipulate my knowledge of what avocados look like by giving me a defeater, and thereby “switching off” my belief that avocados look W, we would thereby “switch off” my competence-grounded knowledge that that’s an avocado. In the chapter, I detail several of these dependences and show why they best explain the truth of the crucial counterfactual mentioned above. Thus, I defuse the Only-Because Argument and defend the possibility of immediate perceptual knowledge of a perceived individual’s kind properties.

In Chapter 2, I consider a second contemporary challenge facing liberalism. *Seemings* play an important role on many theories that recognize immediate justification (perceptual or otherwise) (Huemer 2001, 2005, 2007; Tucker 2010; Bergmann 2013a; Brogaard 2013b; Pace 2010). However, there is little consensus about what seemings are, and a corresponding skepticism about why we should posit the existence of seemings in the first place (Hawthorne & Lasonen-Aarnio ms). In Chapter 2, I develop and defend the epistemic feeling account of seemings. Epistemic feelings are a genus of phenomenally conscious mental states that motivate epistemic performance; familiar examples include the eureka experience (or aha! moment) and the tip of the tongue feeling, or the feeling that one is about to retrieve some piece of information. As part of certain dual process theories of cognition, psychologists have posited an epistemic feeling called the *feeling of rightness* to explain when subjects are likely to simply endorse the output of some Type-1 process in judgement without engaging in Type-2 reasoning (Thompson et al. 2011,

Thompson 2009, Simmons & Nelson 2006). I characterize of the content, phenomenal character, and functional role of the feeling of rightness and argue that central cases of *prima facie* seemings are plausibly understood as cases of feelings of rightness. I also suggest that our central cases of putative seemings cannot be better explained by appealing only to subjects' dispositions to believe.

The epistemic feeling view has a methodological advantage. I develop the view without relying on intuitions about the felicitousness of *seems* expressions, which we might worry can lead to equivocation on *seems*. Moreover, the epistemic feeling view does not depend on holding that the correct account of immediate perceptual justification must appeal to seemings. The primary motivation for positing seemings, on the epistemic feeling view, is the role they play in psychological explanation of subjects' doxastic attitudes, rather than the role they play in explaining linguistic intuitions or justification. (Thus, the epistemic feeling view is compatible with the notion of recognitional competence that I develop in Chapter 1, and with the basic method strategy that I favor in Chapter 4, which can be cashed out in ways that don't appeal to seemings). That the epistemic feeling view is independent of how these controversial issues turn out is, I think, an advantage of the view.

In Chapter 3, I consider another serious contemporary challenge to liberalism: that liberal views of perceptual justification are incompatible or in tension with Bayesian confirmation theory. White (2006) has pointed out that because ordinary propositions about the external world, (e.g., I have a hand), entail the falsity of certain skeptical hypotheses, (e.g., I'm a BIV who merely seems to have a hand), one can show the following:

**Inequality:**  $\text{Cr}(H|E) < \text{Cr}(\sim\text{SK})$ , where:

H: I have a hand.

SK: I am a BIV who merely seems to have a hand.

E: I seem to have a hand.

*Inequality* seems, at least at first glance, to support conservatism, or the view that one has justification to believe external world proposition H partly in virtue of having independent justification (or default entitlement) to believe the denial of skeptical hypotheses, e.g.,  $\sim\text{SK}$  (Wright 2002, 2004, 2007; Davies 2000, 2003).

I first consider the most natural response to this challenge, exemplified by Moretti (2015). This is to point out that the argument assumes that subjects acquire perceptual justification by updating by strict conditionalization on propositions like E—*I seem to have a hand*. However, on a liberal view, acquiring immediate perceptual justification needn't involve adopting credence 1 in a proposition about one's own mental states.

However, the argument can be revised in a way that is more charitable towards the liberal. White (2006) presents cases in which it seems, intuitively, that the rational response to an ordinary perceptual experience, given one's background evidence, is to increase one's credence in a skeptical hypothesis. But if that's right, then it can be shown about those cases that  $\text{Cr}(H|E) < \text{Cr}(\sim\text{SK})$ , and the argument against liberalism can be recovered. I defend liberalism by showing how the revised argument depends on certain assumptions about the rational force of higher-order evidence. Once we recognize the

role that higher-order evidence is playing in driving the crucial intuitions, it becomes clear how White's cases are disanalogous from ordinary cases of perceptual knowledge, and so do not threaten liberalism.

Finally, in Chapter 4, I continue the work of Chapter 1 in broadening the scope of immediate perceptual justification. I consider a natural way of addressing the scope challenge: to hold that the content of perceptual experience delimits the range of beliefs that can be immediately perceptually justified. I argue against this natural answer by pointing to a class of cases in McGrath (2016) and Silins (2011), which I call *mismatch cases*, in which, I suggest, subjects are immediately perceptually justified in believing *p* even though *p* is not among the contents of their perceptual experiences. (By contrast, one use of the more familiar speckled hen case is to show that S can fail to have immediate perceptual justification to believe that *p* even when *p* is indeed among the contents of her perceptual experience.) The gap in content, moreover, does not owe to any general difference in the type of content that belief and experience can have, e.g., the difference between non-conceptual and conceptual content, or iconic and discursive content.

After arguing that the mismatch cases should be characterized in this way, I examine the implications for liberalism. I argue that a version of phenomenal conservatism, according to which S has immediate *prima facie* justification to believe *p* iff it seems to S that *p*, can't explain the possibility of mismatch cases because it draws too tight a connection between the content of perceptual experience and the content of immediately justified perceptual belief. Finally, I point to two liberal strategies for handling mismatch cases. What I call the *permissive evidentialist strategy* accounts for

mismatch cases by adopting a permissive notion of the evidential support relation that must hold between an experience and a belief that it immediately justifies. What I call *the basic method strategy* handles mismatch cases by accounting for immediate perceptual justification in terms of non-inferential competences, methods, and processes. In the final section of the paper, I briefly argue that the basic method strategy is the more promising solution.

## Chapter 1

### Recognitional Competence and Knowing What Things Look Like

When I see an avocado in the grocery store, I can know that it's an avocado even if it's unlabeled. Moreover, the following counterfactual seems true about me: If I didn't know what avocados look like, I wouldn't know that that's an avocado. These considerations motivate the thought that my perceptual knowledge that that's an avocado depends in some way on my knowledge of what avocados look like.

McGrath's (2017) *Only-Because Argument* purports to show that my perceptual knowledge that that's an avocado is *based on* my knowledge of what avocados look like, which, McGrath argues, is propositional knowledge. If the argument succeeds, it shows that perceptual knowledge that an ordinary object has a certain kind property, e.g., my knowledge that that's an avocado, is mediate knowledge because it is based on a piece of propositional knowledge, e.g., my knowledge of what avocados look like. This result would seriously restrict the scope of immediate perceptual knowledge.

In this paper, I defend the possibility of immediate perceptual knowledge of an ordinary object's kind properties. I suggest that such immediate perceptual knowledge is grounded in recognitional competence. I argue that we can appeal to recognitional competence to explain why my perceptual knowledge that that's an avocado depends counterfactually on my knowledge of what avocados look like, even if we assume that McGrath is right that to know what avocados look like is to possess propositional knowledge. Indeed, the recognitional competence account better explains this

counterfactual dependence than does the view that my knowledge that that's an avocado is mediate knowledge.

Here's the plan for the paper. In the first section, I present the Only-Because Argument. In Sections 2 and 3, I argue that a premise in the argument is false by developing an account of recognitional competence. Sections 4-6 are devoted to arguing that the recognitional competence account does a better job of explaining the crucial counterfactual dependence than the mediate knowledge account.

## 1. The Only-Because Argument

### *1.1 Immediate vs. Mediate Knowledge*

What's at stake in this paper is the possibility of a certain sort of immediate perceptual knowledge. Following McGrath (2017, 5) and Alston (1983), we can distinguish immediate and mediate knowledge in the following way:

**Immediate Knowledge:** S has immediate knowledge that p just when S's belief that p is based on an immediate epistemizer S has for p.

**Mediate Knowledge:** S has mediate knowledge that p just when S's belief that p is based on a mediate epistemizer S has for p.

An epistemizer for  $p$  is a factor that, when a subject possesses it, helps position her to know that  $p$ . “Epistemizer” is a generic term; what kind of factors one counts as epistemizers will depend on one’s background theory of knowledge and justification. On most views, a piece of propositional knowledge can count as an epistemizer. My propositional knowledge that orange trees are tropical plants and that Alaska isn’t a tropical region are epistemizers that help put me in a position to know that orange trees don’t grow in Alaska. I take advantage of my epistemic position by basing my belief that orange trees don’t grow in Alaska on these epistemizers. Other epistemizers might be abilities, competences, skills, or reliable processes. I might, for example, have available a suite of reliable processes through which I can recognize things as avocados when I see them. On certain views of knowledge and justification, having available these processes constitutes an epistemizer for the proposition *that’s an avocado*. I take advantage of my epistemic position, in this case, by using those processes in coming to believe that *that’s an avocado*.

Whether an epistemizer is mediate or immediate depends on whether it includes a piece of propositional knowledge.

**Immediate Epistemizer:** “An epistemizer,  $E$ , one has for  $p$  is immediate iff: (i) no factual knowledge one has appears in  $E$ , and (ii) one has no mediate epistemizer,  $E^*$ , for  $p$ , which is such that  $E$  positions one to know that  $p$  only in virtue of  $E^*$  positioning one to know that  $p$ ” (McGrath 2017, 6).

**Mediate Epistemizer:** An epistemizer E that S has for p is mediate iff it is not immediate.

Factual knowledge here just is propositional knowledge. So if what positions me to know that p includes propositional knowledge, then my knowledge that p is mediate. My knowledge that orange trees don't grow in Alaska is mediate because it's based in part on other propositional knowledge, e.g., my knowledge that orange trees are tropical plants.

What does it mean for propositional knowledge to “appear in” or “be included in” an epistemizer? One way for propositional knowledge to appear in an epistemizer is for that epistemizer to be identical with a piece of propositional knowledge. Another way is for the epistemizer to be a set of factors, at least one of which is identical to a piece of propositional knowledge. By contrast, the second clause in the definition of an immediate epistemizer highlights cases in which a piece of propositional knowledge doesn't *appear in* an epistemizer, but rather grounds an epistemizer. McGrath's example here is *understanding chess* (2017, 5-6). Understanding chess might be an epistemizer that puts me in a position to know that rooks can't move diagonally. I might understand chess partially in virtue of possessing some propositional knowledge, e.g., that rooks can only move horizontally and vertically. Moreover, it might seem that my understanding chess positions me to know that rooks can't move diagonally only because I understand chess partly in virtue of possessing those other pieces of propositional knowledge. Plausibly, it is my propositional knowledge that is doing the real work in explaining how I know that rooks can't move diagonally. Thus, we'll want to say that my knowledge that rooks can't move diagonally is mediate.

## 1.2 The Only-Because Argument

Now we are well-positioned to present the *Only-Because Argument*. It starts with the following case, adapted from McGrath (2017, 8).

**Supermarket:** My friend Dmitri and I work together in an eco-friendly supermarket that never labels its produce. We both have seen what are in fact avocados many times and can respond differentially in behavior to the way they look, e.g., by sorting avocados from mangos and papayas. Moreover, we both know certain facts about avocados, e.g., that they are used to make guacamole, that California grows a lot of them, etc. However, neither of us has had occasion to learn that the things that we regularly sort, which look a certain way, are avocados. One day, while Dmitri isn't around, the manager refers to the avocados I'm sorting as avocados. It's natural to say that I now know what avocados look like, while Dmitri does not. Moreover, when I see the unlabeled avocados, I can know that they are avocados. Dmitri, on the other hand, does not know that they are avocados. Let's suppose that he suspends judgment both on the proposition that they are avocados, and on any proposition about what avocados look like.

The first thing to notice about the case is that there is a special sense in which I know what avocados look like, while Dmitri does not. Dmitri might be said to know (*de re*) what avocados look like, because he can respond differentially to avocados by sorting

them reliably from other fruit. However, Dmitri does not know (*de dicto*) what avocados look like because he doesn't know that the things he regularly sorts, which look a certain way, are avocados, i.e., the very same things that he believes are in guacamole, etc. On the other hand, I do know (*de dicto*) what avocados look like. It is this *de dicto* reading of knowing what avocados look like that is relevant to the argument (McGrath ms, 2).

The second thing to note about the case is that the following counterfactual seems true about it: If I didn't know what avocados look like, I wouldn't know that those are avocados. To put things differently, I know that those are avocados *only because* I know what avocados look like. This dependence supports the thought that any epistemizer I have for the belief that those are avocados includes or is grounded in my knowledge of what avocados look like. If I had some other epistemizer for the belief that those are avocados, (e.g., if I had seen a label), then I would still be in a position to know that they are avocados, even if I didn't know what avocados look like. McGrath lays out this reasoning as follows, replacing "avocados" with Fs to make the argument general.

### **The Only-Because Argument**

1. I know that those are Fs, and I know what Fs look like.
2. If I have an epistemizer for those are Fs that doesn't include my knowledge of what Fs look like, then if I didn't know what Fs look like, I'd still be in a position to know that those are Fs.
3. However, if I didn't know what Fs look like, I wouldn't be in a position to know that those are Fs.

4. Every epistemizer I have for those are Fs includes my knowledge of what Fs look like. (From 2 and 3).
5. Knowledge of what Fs look like is propositional knowledge.
6. My knowledge that those are Fs is mediate knowledge. (From 1, 4, and 5) (McGrath 2017, 10-11).

Premise 1 is a plausible stipulation about the supermarket case and so should be unproblematic. Premises 2 and 3, at least at first glance, seem to be true intuitively. However, Premise 5—that knowledge of what Fs look like is propositional knowledge—is controversial.<sup>2</sup> McGrath supports Premise 5 in part by developing an intellectualist account of knowing what Fs look like, along the following lines.

**Simple Intellectualist Account:** S knows what Fs look like just when S has propositional knowledge that Fs look W.<sup>3</sup>

My response to the Only-Because Argument requires that we unpack this view a bit. First, W is a special concept or way of thinking that refers to a “looks property” *w*. A looks property, in this context, is not some property of mental states; it is not identical with the phenomenal character of normal perceivers’ experiences as of Fs. Rather, a

---

<sup>2</sup> Shieber (2017), for example, challenges Premise 5 by arguing that to know what Fs look like is to possess a competence or ability, rather than to have some piece of propositional knowledge. According to Shieber, S knows what Fs look like just when she has “the ability visually to distinguish Fs from non-W-looking things” (4). McGrath (ms) responds in part by suggesting that, while Shieber’s account may be plausible for knowing *de re* what Fs look like, it is not a plausible account of knowing *de dicto* what Fs look like. Dmitri, after all, can visually distinguish avocados from papayas and mangos, but he does not know *de dicto* what avocados look like.

<sup>3</sup> See also McGrath (2018) for further development of the intellectualist view.

looks property is a property (or set of properties) of external world objects (McGrath 2018, 24). In this usage, an apple—not an experience of an apple—has a looks property. There are a couple options for understanding looks properties; we can remain neutral on these. Looks properties could be (1) sets of ordinary perceptible properties, e.g., color, shape, size, texture, etc., (2) sets of viewpoint-relative properties that are partly determined by the perceiver's environment, e.g., Schellenberg's (2008) situation-dependent properties, (3) a combination of 1 and 2. The generic—Fs look W—asserts that instances of F typically or characteristically have looks property w, (in normal conditions). Thus, since there is not always some way of looking that members of a kind have in common, only some kinds will have associated looks properties. Moreover, there will likely be many cases in which it is vague whether some object falls in the extension of the concept W.

We can also remain neutral among the several options that McGrath offers for how to understand W, the concept that refers to looks property w. When we think of something as being W, we are thinking of that thing as having looks property w. W might then be what Papineau (2006) calls a perceptual concept, or a concept “involving stored sensory templates established when we perceive an object or feature and which are reactivated upon subsequent perceptual matches as well as in imagination” (McGrath 2017, 16). Alternatively, W might be a recognitional concept. For S to possess the concept W, on this type of account, S must possess a skill, ability, or competence to recognize the referent of W—i.e., looks property w—by reliably discriminating w from other looks properties in behavior and belief. Dmitri, for example, might possess a concept W of the look of an avocado, even though he does not have propositional

knowledge that avocados look W. For our purposes, it won't matter much exactly how we think of W. I will simply assume that Premise 5 is true. There is a way to resist the argument even assuming intellectualism is true.

According to intellectualism, in the supermarket case, I come to know what avocados look like because I acquire the propositional knowledge that avocados look W. Because I now know that avocados look W, I'm in a position to competently perform the following tacit inductive inference and so come to know its conclusion:

### **Tacit Inference**

1. Avocados look W.
2. This thing looks W.
3. So this thing looks like an avocado.
4. If this thing looks like an avocado, then (probably) it is an avocado.
5. This thing is an avocado (probably).

Thus, the intellectualist has a neat explanation for why my knowledge that those are avocados depends counterfactually on my knowing what avocados look like.<sup>4</sup> Namely, if I didn't know what avocados look like, I wouldn't know that avocados look W, and so I wouldn't be able to competently perform this tacit inference and come to know its conclusion.

---

<sup>4</sup> Tacit Inference represents a natural explanation of the subject's justification in the supermarket case for the simple intellectualist who endorses the Only-Because Argument. However, strictly speaking, the simple intellectualist account is just a view about knowledge of what Fs look like. Thus, the simple intellectualist doesn't have to give this account of the supermarket case. Indeed, the simple intellectualist account is compatible with holding that the subject's belief in the supermarket case is immediately epistemized.

If the Only-Because Argument is sound, it shows that perceptual knowledge that attributes a kind property to an ordinary object must be mediate knowledge, because it must be based on my knowledge of what that kind looks like. For McGrath, the relevant kind properties form a large class: They include:

[A]rtifactual kinds at all levels of genus and species (e.g., musical instrument, oboe, Fox Renard oboe), and natural kinds at all levels (e.g., animal, dog, Scottish Terrier), and well as just about any sortals we ordinarily apply to categorize objects, again at all levels (e.g., rock formation, mountain). (2017, 111)

I suspect, moreover, that if the argument is sound, it would generalize to so-called “low-level” observable properties such as color, shape, size, distance, pitch, loudness, bitterness, etc. For example, it seems that if I didn’t know what pentagons or chartreuse look like, then I wouldn’t know that *that’s a chartreuse pentagon* when I see one, (in cases, analogous to the supermarket case, in which I lack some non-perceptual epistemizer for that proposition). Likewise, if I didn’t know what bitterness tastes like, I wouldn’t know that *that’s bitter* when I taste it, (again, in cases where I lack some non-perceptual epistemizer). So plausibly, the argument threatens to undermine, not just immediate perceptual knowledge of a perceived object’s kind properties, but also immediate perceptual knowledge of a perceived object’s so-called “low-level” properties. Defenders of immediate perceptual knowledge, then, should be particularly eager to respond to the argument.

## 2. Recognitional Competence

I will challenge Premise 2 of the Only-Because Argument. By challenging Premise 2, we thereby block the crucial move to 4: that every epistemizer I possess for

the belief that *those are avocados* includes my knowledge of what avocados look like.

Premise 2 is a material conditional with a counterfactual as its consequent: If I have an epistemizer for those are Fs that doesn't include my knowledge of what Fs look like, then if I didn't know what Fs look like, I'd still be in a position to know that those are Fs. I'll argue that the antecedent of the conditional can be true while the consequent is false. That is, I will explain why it can be the case that:

- A. Subject S (from the supermarket case) has an epistemizer for those are Fs that doesn't include her knowledge of what Fs look like.
- B. It's not the case that: if S didn't know what Fs look like, she would still be in a position to know that those are Fs.

First, I will sketch a way in which A can be true, i.e., how a subject like the one from the supermarket case can have an epistemizer for the proposition that *those are Fs* that doesn't include her knowledge of what Fs look like, which we're assuming is propositional knowledge that Fs are W. I'll argue that the relevant epistemizer is a recognitional competence, and explain why that epistemizer doesn't include the propositional knowledge that Fs are W. That will be the task of this section. The task in the rest of the paper will be to explain how B can be true—how my knowledge that those are Fs can depend counterfactually on my propositional knowledge that Fs are W—even though I have an immediate epistemizer for the belief that those are Fs, in the form of my recognitional competence.

Recall again the supermarket case. When I learn that the things I've been sorting, which look a certain way, are avocados, I acquire a recognitional competence, or a certain sort of reliable disposition. There are a variety of possible dispositional structures that we could use.<sup>5</sup> I'll develop one in detail, but my argument should go through on alternative accounts of the structure of recognitional competence. Here's a first pass at a characterization of the relevant disposition, which we'll go on to revise:

**Recognitional Competence Account (RCA):** S has a competence to recognize Fs iff: S is disposed to believe truly that x is F (absent defeaters) if she perceives that x is G (under relevant conditions).

Let's unpack RCA before clarifying how it relates to the Only-Because Argument. I intend RCA to be ecumenical about several issues in the philosophy of perception. First, for concreteness, I will often present RCA by appeal to the representational content of experience, but we could recast RCA to be compatible with an adverbialist or naive realist account of perception, e.g., by appealing to S's sensing G-ly, or to S's being perceptually related to an instance of G-ness. Second, within the representationalist camp, RCA does not take a stand on whether the contents of perceptual experience are general, e.g.,  $\exists x Gx$ , or particular, e.g., that's G. We should read "x is G" as being neutral between

---

<sup>5</sup> For example, Sosa (2003; 2015, 202-205) argues that the problem of the speckled hen for classical foundationalism must be solved by appeal to recognitional competence. Pace (2010) extends this argument by detailing the structure of recognitional competence, which he understands roughly as a disposition to form reliably true beliefs that a is  $\phi$ , (where  $\phi$  is a concept that refers to a perceptible property, *Gness*), when S tokens a perceptual experience that represents that a is G (429). By contrast, Millar (2011, 334) gives a knowledge-first account of recognitional abilities according to which exercising a recognitional ability requires not merely that one form a true belief, but moreover that one acquire perceptual knowledge.

these two options. Finally, while *x* is meant to be an external world individual, rather than a private mental object, RCA is neutral on exactly what sort of external world individuals we can perceive as being *G* and recognize as being *F*. We can leave open the possibility that *x* is (1) an event, e.g., a car horn blaring, (2) an ordinary object, e.g., a cat, (3) a non-ordinary but still perceptible object, e.g., a rainbow, hole, or shadow, (4) a sound or smell, so long as these are understood as external world individuals. RCA can cover cases in which we perceptually recognize any of these sorts of individuals as being *F*, i.e., as having a kind property.

Next, we should understand *G*-ness as a property (or set of properties) of external world individuals. I'm thinking of *G*-ness as a property (or set of properties) such that an external world individual looks/sounds/smells/feels/tastes like an *F* (or some combination of these) in virtue of instantiating *G*-ness. The relevant sense of "looks like" is Jackson's (1977) comparative sense. In the comparative sense, an individual looks like an *F* to *S* in circumstances *C* when it looks the way that *F*s look to normal perceivers in *C* (Jackson 1977, 32). Note, then, that *G*-ness is a stand-in for a wide-range of properties (and sets of properties), since there needn't be any single property (or set of properties) in virtue of which an individual looks/sounds/smells/etc. like an *F* to normal perceivers in *C*. For example, an object could look like a tree in virtue of having willow-tree-look or pine-tree-look, which are plausibly incompatible with each other. So if *S* has an experience representing that *x* is *G* in circumstances *C*, she has an experience that represents an individual as having *some* property (or set of properties) in virtue of which *x* looks/sounds/smells/etc. like an *F* to normal perceivers in *C*.<sup>6</sup>

---

<sup>6</sup> One assumption I'm making here is that experiences themselves don't represent individuals as instantiating kind properties, contra Siegel (2010) and Bayne (2009). Defending this assumption

It is only S's disposition to form reliably true categorization beliefs in a certain range of conditions that is relevant to whether S possesses a recognitional competence. For example, that S fails to form true categorization beliefs when her experiences are illusory or hallucinatory shouldn't count against her possessing a recognitional competence. We should distinguish competences to discriminate and single out individuals in perception from recognitional competences, whose exercise involves deploying concepts in belief. So I suggest that we evaluate S's competence to recognize Fs by assessing the reliability of her categorization beliefs in cases where she not only has an experience representing that x is G, but where that experience constitutes successful perception.

Moreover, we should clarify how RCA relates to attention. Suppose we assume that S's perceptual experience can represent features that are attended to only a very low degree or not at all. In that case, S might fail to believe that x is F, even though she has an experience representing that x is G, because she fails to attend sufficiently to x's G-ness. Should this count against S's possessing a competence to recognize F? I don't think so. I suggest that we evaluate S's competence to recognize Fs by assessing her dispositions to respond to experiences in which she attends to a high degree to x's G-ness. It is true that, through perceptual learning, a subject might develop dispositions to attend to features of individuals in ways that facilitate recognition. I suggest we think of these attentional dispositions as putting subjects *in a position to* exercise recognitional competences. So for example, I might have acquired a disposition, through perceptual learning, to attend to the shape and number of the leaves in determining whether a plant is poison ivy. By

---

in a responsible way would take us too far afield. For an argument that the content of experience is "thin," see Byrne's contribution to Byrne & Siegel (2017).

exercising this attentional disposition, I might come to have an experience representing that a plant has a particular shape, color, and number of leaves, in which I attend to these features. That attentive experience puts me in a position to exercise my competence to recognize poison ivy by coming to believe truly that that's poison ivy in the absence of defeaters.

Finally, RCA is simplified in an important respect. Compatibly with satisfying RCA, S might also have a disposition to believe that x is F in response to experiences that represent x as H, where *being H* is incompatible with *being G*, and in virtue of *being H* x does not look like an F. For example, in the supermarket case, by the lights of RCA, I can have a competence to recognize avocados even if I am disposed to believe that *that's an avocado* when I see, not just avocados, but also papayas. To avoid this problem, we can revise the account as follows:

**Recognitional Competence:** S has a competence to recognize Fs iff:

1. S is disposed to believe truly that x is F (absent defeaters) if she perceives that x is G (under relevant conditions).
2. It's not the case that S is disposed to believe that x is F if she perceives that x is H, (all other things being equal), where  $\forall x(Gx \rightarrow \sim Hx)$ .

So if I'm disposed to respond to visual experiences as of papayas by believing that they are avocados, I violate 2 and so don't count as possessing a competence to recognize

avocados. Perhaps surprisingly, this is all compatible with saying that there can be many, incompatible ways of looking like an F. For example, something can look like a tree in virtue of having willow-tree-look or pine-tree look. Thus, a subject can exercise a competence to recognize trees by responding in the appropriate way to perceiving either look; both perceptions would count as a perception that “x is G,” since an individual can count as looking like an F by having a wide range of properties (or sets of properties).<sup>7</sup>

Now we are in a position to connect the account of recognitional competence with the Only-Because Argument. Recall again that our task in this section is to explain how A can be true:

A. A subject S has an epistemizer for those are Fs that doesn’t include her knowledge of what Fs look like, i.e., her propositional knowledge that Fs look W.

The proposal is that we understand the relevant epistemizer as a competence to recognize Fs.

Although often we acquire a competence to recognize Fs at the same time as we acquire propositional knowledge that Fs look/sound/etc. W, the two are distinct. It’s conceptually and metaphysically possible for a subject to possess propositional knowledge that Fs look W, without possessing a competence to recognize Fs, as in the following case adapted from McGrath (2017, 26-7):

---

<sup>7</sup> Thanks to Frankie Egan for pointing out the need for clarification here.

**Ringers:** Suppose that two species of butterfly in the Amazon, Es and Fs, look exactly the same to normal human perceivers. Thus, Brent, the entomologist, lacks a competence to recognize Fs. He doesn't have a disposition to believe truly that x is an F in response to appropriately attentive experiences representing that x is G. After all, Fs and Es look exactly alike. Nevertheless, Brent still knows what Fs look like. That is, Brent knows that Fs are W. (The same is true for Es. Brent can know that Es look W, though he lacks a competence to recognize Es.)

The upshot is that it is conceptually and metaphysically possible to know what Fs look like, (i.e., to know that Fs are W), without possessing a competence to recognize Fs. (The possibility of ringers is why we cannot simply identify knowing what Fs look like with possessing a competence to recognize Fs, although that identification seems initially tempting.)

It's also conceptually and metaphysically possible for S to possess a competence to recognize Fs even though she doesn't count as knowing what Fs look like, as in the following case:

**Gettier Case:** Imagine that you're the subject from the supermarket case. You have seen what are in fact avocados many times and can respond differentially in behavior to the way they look, e.g., by sorting avocados from mangos and papayas. But you don't yet know that these things, which you regularly sort, are avocados. One day, as you're sorting the avocados, your friend tells you that these are avocados. Suppose that she is a highly reliable testifier in general. Thus, you

come to believe truly and justifiably that avocados are W. However, your friend is trying to play a trick on you; she believes that the things you are sorting are really guavas and is trying to give you a false belief. But fortunately for you, she has mixed up guavas and avocados. Thus, you're Gettiered, and so don't have propositional knowledge that avocados are W. Still, you might thereby acquire a recognitional competence for avocados. You are now disposed to reliably form true categorization beliefs, e.g., *that's an avocado*, in the absence of defeaters, when you see objects with the relevant look.

Thus, it is possible to possess a recognitional competence for Fs even if you don't possess propositional knowledge that Fs are W. The upshot is that these two epistemizers are not type-identical. Possessing a competence to recognize Fs is neither necessary nor sufficient for knowing what Fs look like.

### **3. Explaining Counterfactual Dependence**

In this section, I'll propose an explanation for the crucial counterfactual compatible with RCA. In the rest of the paper, I'll argue that this explanation enjoys an advantage over the most plausible competing explanation, i.e., that my knowledge that that's an avocado is inferentially based on propositional knowledge that avocados look W.

Here's the crucial counterfactual again: If S didn't know what avocados look like, then S wouldn't know that that's an avocado. Importantly, S is the subject of the

supermarket case; the interpretation of the counterfactual is in part fixed by the case. The thought is that if S didn't know what avocados look like, she'd be relevantly like Dmitri, and so wouldn't know that *that's an avocado* upon seeing one.

Why, according to the RCA, is the counterfactual true? It will be difficult to answer this question in any detail without appealing to some view of the semantics of counterfactuals. However, the correct semantics for counterfactuals is a controversial and difficult issue that would take us too far afield. I suspect, moreover, that the solution I'll propose doesn't depend too closely on any one account. My strategy, then, will be to illustrate the RCA's explanation of the crucial counterfactual using Lewis's (1973) semantics for counterfactuals, including his (1979) account of the determinants of the similarity relation. I'll then briefly indicate why I think the explanation will also work with premise semantics for counterfactuals (Kratzer 1989, 2002; Veltman 2005). In the rest of the paper, I'll argue that giving an interventionist (or causal model) interpretation of the counterfactual actually supports the RCA.

According to a Lewisian (1973) similarity account, a counterfactual  $\phi > \psi$  is true at  $w$  iff all of the worlds most similar to  $w$  at which  $\phi$  is true are worlds in which  $\psi$  is true. The relevant question for us, then, is what are the closest possible worlds in which S, the subject of the supermarket case, doesn't know what avocados look like. There are many ways to fail to possess the propositional knowledge that avocados are W. It could be false that avocados are W; S could fail to believe or justifiably believe that avocados are W. Or S's justified true belief that avocados are W could be Gettiered. Suppose Lewis (1979) is right that we preserve similarity by avoiding diverse, widespread violations of law, and less importantly, by avoiding divergence in particular matters of fact. It strikes

me that the closest worlds in which I fail to know what avocados look like are the ones in which I fail to believe that avocados look W. (I would then be relevantly like Dmitri, who also fails to believe that avocados look W.) One might think it only takes a small, localized violation of law to alter my mental states in this way, but it would take a much more widespread violation of law and divergence of particular fact to change the way avocados in fact look, or to ensure that my belief is Gettiered or fails to be doxastically justified. So let's say that the Lewisian similarity relation selects those worlds in which I don't know what avocados look like because I fail to believe that they look W (either because I disbelieve this, suspend judgment on it, or fail to possess any doxastic attitude towards it).

What else would be true in a world in which I failed to believe that avocados are W? Plausibly, holding fixed that I'm rational, I would also fail to possess the disposition that in part constitutes the competence to recognize avocados. That is, I would fail to be disposed, in the absence of defeaters, to transition from experiences that represent W-looking things to beliefs that *that's an F*. The motivation for this thought is the familiar idea that a rational subject who instantiates a particular propositional attitude is disposed to behave and reason in certain ways, depending on her other mental states. (I intend this to be a minimal claim about rational subjects that is compatible with dispositionalist, functionalist, and radical interpretationist accounts of propositional attitudes, among others.)

Let's consider three possibilities: Rational subject S fails to believe that Fs are W because either (1) S disbelieves that Fs are W, (2) S suspends on whether Fs are W, or (3) S has no doxastic attitude on whether Fs are W. In the first two cases, it is not part of the

dispositional profile of the respective doxastic attitudes that the subject possess the disposition that underlies recognitional competence, i.e., the disposition to transition from experiences as of W-looking things to the belief that that's an F. And the same will be true for the third case: If one takes no doxastic attitude towards the relevant proposition, one won't possess any particular dispositional response to experiences as of W-looking things, all other things equal.

One might object in the following way. There are cases in which subjects who disbelieve that Fs are W are still disposed to transition from experiences that represent that x is W to the belief that x is an F.<sup>8</sup> Suppose, for example, that W is a concept of the look of a ponderosa pine. Rational subject S might truly disbelieve the generic—trees are W—since trees do not characteristically have the ponderosa-pine-look. Still, she might have the disposition to transition from experiences as of W-looking things to the belief that *that's a tree*. So in at least some cases, subjects who disbelieve that Fs are W might still possess a disposition that could constitute a recognitional competence.

But suppose we focus on the particular case that fixes the interpretation of the counterfactual. Here the relevant attitude is disbelieving that (or suspending on whether) avocados are W, where W is in fact the look of an avocado. This case is disanalogous from the ponderosa pine case: We are not considering the look of one of the determinates of F. So it seems safe to say that *this* subject, who disbelieves that (or suspends on whether) avocados are W, would lack the disposition underlying the competence to recognize avocados.

---

<sup>8</sup> Thanks to Matt McGrath for pointing this out.

The suggestion is that, in all the closest worlds in which I don't know what avocados look like, I also fail to possess the disposition that partially constitutes a competence to recognize Fs. Namely, I fail to be disposed to believe that that's an F in the absence of defeaters in response to the right sorts of perceptual experience. (Possessing that disposition is not sufficient to possess a recognitional competence because the disposition might fail to be reliable.) According to RCA, I wouldn't then know that *that's an avocado* when I see one, because I lack the recognitional competence. Thus, we have an explanation of the crucial counterfactual that is compatible with RCA.

The same general type of explanation will be available on alternative accounts of the semantics of counterfactuals. Consider, for example, Veltman's (2005) premise semantics. Explaining this view in full detail would take us too far afield, but it is worth noting that it underwrites the same sort of explanation of the crucial counterfactual. Briefly, here is why. On Veltman's view, context supplies a set of worlds and generalizations that are relevant to interpreting a counterfactual. Evaluating a counterfactual requires that one *retract* the fact that the antecedent is false from the worlds supplied by the context. This retraction involves, first, finding the *basis* of each world *w*, or the smallest set of facts at *w* that, given the generalizations supplied the context, determine the rest of the facts about that world. One then retracts the antecedent from the basis of each world *w* by changing the fewest propositions in the basis such that the basis no longer forces the antecedent to be true, given the generalizations supplied by the context. The relevance of this for us is this: It seems plausible that, in evaluating the crucial counterfactual, retraction of the antecedent would proceed by retracting the fact

that S believes that Fs are W, since this could easily represent the minimal change to the basis that allows the antecedent of the counterfactual to be true. However, as we've seen, worlds in which S fails to believe that Fs are W are also worlds in which she lacks a recognitional competence and so fails to know, by exercising that competence, that *that's an F*. This means that RCA's explanation of the crucial counterfactual is available on other accounts of the semantics of counterfactuals besides Lewis's.

Thus, the recognitional competence account can explain why the crucial counterfactual is true. Recognitional competence and propositional knowledge that Fs are W are closely related to each other, since a rational subject who possesses either one will have a certain classificatory disposition. However, as the Ringers and Gettier cases show, neither is necessary for the other. Moreover, our explanation of the crucial counterfactual does not require that in order for S's recognitional competence to yield knowledge, S must also possess propositional knowledge that Fs are W. So neither S's possession of a recognitional competence, nor that competence's status as an epistemizer, requires that S know that Fs are W. This concludes the preliminary argument for A and B:

- A. Subject S (from the supermarket case) has an epistemizer for those are Fs that doesn't include her knowledge of what Fs look like.
- B. It's not the case that: If S didn't know what Fs look like, she would still be in a position to know that those are Fs.

If A and B are true, then Premise 2 is false. The task in the rest of the paper will be to bolster the preliminary argument. I'll suggest that we cannot better explain B by holding that my knowledge that those are avocados is inferentially based on my knowledge that avocados look W.

#### **4. Interventionist Counterfactuals and Metaphysical Dependence**

An immediate worry with the foregoing account is that it depends too closely on certain views of the semantics of counterfactuals. In the rest of the paper, I'll attempt to address the worry by considering an interventionist interpretation of the counterfactual. According to interventionist accounts, when a counterfactual expresses a causal relationship between events or states of affairs in the antecedent and consequent, the counterfactual is made true by a fact about what would happen under a certain sort of intervention. An intervention on a variable X with respect to a variable Y is a "surgical change" to the value of X such that if Y would thereby change, it would be due entirely to the causal influence of X on Y.<sup>9</sup> So a counterfactual that expresses a causal relationship between the antecedent and consequent is made true by a fact about what would happen under an intervention on the antecedent with respect to the consequent.

An interventionist interpretation of our crucial counterfactual would, then, be favorable towards the competing view: that my knowledge that that's an avocado is inferentially based on (and so presumably caused by) my propositional knowledge of what avocados look like. On this picture, because this causal relationship obtains, the

---

<sup>9</sup> For interventionist accounts of causation, see Woodward (2003), Woodward & Hitchcock (2003), and Pearl (2009). For interventionist semantics for counterfactuals, see Briggs (2012).

crucial counterfactual is made true by a fact about what would happen under an intervention on the antecedent with respect to the consequent: If we were to intervene surgically to “switch off” my knowledge of what avocados look like, we would thereby “switch off” my knowledge that those are avocados.

In the rest of the paper, I will argue that my proposed explanation of the crucial counterfactual enjoys an advantage over the interventionist explanation. Namely, I’ll argue that it is *metaphysically impossible* to perform the relevant intervention. On an interventionist picture, the impossibility of the relevant intervention means that the counterfactual is not made true by a causal relationship between the variables. I’ll argue that the relevant intervention is metaphysically impossible by considering several of the most plausible putative interventions and showing that each of these is impossible. The RCA can explain the crucial counterfactual without positing this causal relationship, and so it is consistent with the impossibility of the relevant intervention. This, I’ll suggest, is a significant advantage of the recognitional competence account.

To present the argument against the interventionist explanation of the counterfactual, we’ll need to briefly explain (1) what an intervention is, and (2) how we should understand interventions in cases where there exist metaphysical dependence relations, (e.g., supervenience, grounding, identity, definitional dependence), among variables. As we’ve seen, an intervention on variable X with respect to Y is a surgical change on X such that, if Y were thereby to change, it would be due entirely to the causal influence of X. Using a simplified version of Woodward (2003, 2016) and Woodward & Hitchcock (2003), we’ll say that interventions are surgical in the sense that they meet the following necessary conditions:

**Intervention:** I is an intervention on X with respect to Y only if:

1. “I [is] the only cause of X” (Woodward 2016, 20).
2. “I must not directly [change] Y via a route that does not go through X”  
(Woodward 2016, 20).
3. “I leaves the values taken by any causes of Y except those that are on the directed path from I to X to Y (should this exist) unchanged” (Woodward 2016, 20).

Let’s briefly explain each one. First, an intervention sets a variable to a certain value so that the value of the variable no longer depends on its actual causes, but entirely on the intervention. If we’re trying to determine the causal influence of a barometer pressure reading on whether it rains, for example, an intervention on the barometer reading breaks the connection between the reading and its actual causes, e.g., low atmospheric pressure. So the first necessary condition on an intervention helps isolate the causal influence of X on Y by controlling for potential common causes of X and Y.

According to the second and third necessary conditions, an intervention must not directly change the value of Y, or change the values of any causes of Y other than those that fall on the direct causal route from I to X to Y. Conditions 2 and 3 help isolate the causal influence of X by controlling for the confounding effects of other causes of Y. Suppose that X, in this case, is level of inflammation, Y is level of reported pain, and I is administration of a drug. Even if the drug is effective at reducing inflammation, it might

also cause lower reported pain through the placebo effect. So simply administering the drug does not count as an intervention on X with respect to Y, because I violates either Condition 2 or 3: it either directly causes a change in Y, or changes values of causes of Y that do not fall on the causal route from I to X to Y.

To understand the crucial counterfactual, we'll also need to consider what happens when we intervene on a variable that bears, not just causal relations to other variables, but also metaphysical dependence relations such as supervenience, grounding, identity, definitional dependence, etc. We can sharpen this notion of metaphysical dependence by holding that variables in a set bear metaphysical dependence relations to one another when the set violates independent fixability.

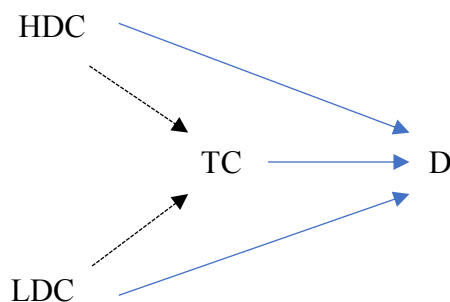
**Independent Fixability:** “(A) set of variables V satisfies independent fixability of values if and only if for each value it is possible for a variable to take individually, it is possible (that is, ‘possible’ in terms of their assumed definitional, logical, mathematical, mereological or supervenience relations) to set the variable to that value via an intervention, concurrently with each of the other variables in V also being set to any of its individually possible values by independent interventions” (Woodward 2015, 316).

Consider, for example, the following simplified case of a causal model that violates independent fixability from Woodward (2015).<sup>10</sup> Whether a subject gets heart disease depends causally on her total cholesterol. Total cholesterol is defined as the sum of high-

---

<sup>10</sup> For other examples that illustrate the problem, see Rescorla (2014).

density and low-density cholesterol (good and bad cholesterol, respectively). So  $TC = HDC + LDC$ . We can represent the situation in the following way, where solid blue arrows represent causal influence and dotted black arrows represent definitional dependence.



Suppose now that we try to intervene on LDC with respect to D. Given how TC is defined, we cannot change LDC while holding fixed both HDC and TC. However, to count as an intervention, I must hold fixed the values of any causes of D that are not on the causal route from I to LDC to D. The upshot is that we cannot evaluate the causal influence of LDC on D in the standard interventionist way, given that our model includes variables that depend metaphysically on one another.

Woodward proposes to solve the problem in the following way. When we are dealing with models that include variables that metaphysically depend on each other, we treat an intervention on one variable as “tantamount to an intervention” on other variables that metaphysically depend on the first. (2015, 331). For example, if we intervene on LDC to change its value while holding HDC fixed, we treat the intervention as automatically changing the value of TC in the appropriate way. We can think of this as a “single intervention” on LDC and TC (2015, 331). In assessing causal influence, we

consider what would happen only under interventions that respect relationships of metaphysical dependence among variables.<sup>11</sup> We're now in a position to explain why all this matters for the Only-Because Argument.

## 5. Intervening on Knowledge Via Belief

In Section 2, I proposed an explanation for the crucial counterfactual—if I didn't know what avocados look like, then I wouldn't know that those are avocados. On a competing explanation, the counterfactual is true because my knowledge that those are avocados is inferred from (and so causally depends on) my knowledge of what avocados look like. If this competing explanation is right, we should expect there to be an interventionist reading of the crucial counterfactual; if we were to turn off the one piece of knowledge through an intervention, we would thereby turn off the other piece of knowledge. In this section and the next, I'll argue that there is no such interventionist reading available because it is metaphysically impossible to perform the relevant intervention. And on an interventionist picture, this means that the counterfactual is not made true by a causal relationship. I'll argue for this by considering some of the most plausible putative interventions and arguing that they fail to satisfy the necessary conditions on interventions. I'll call a merely putative intervention a "manipulation." This

---

<sup>11</sup> Woodward (2015, 330) acknowledges that the revision opens up the possibility of double-counting: The causal influence of X on Y cannot simply be identified with what would happen to Y under a suitable intervention on X, for some of the change in Y may owe to changes in variables that metaphysically depend on X, and so must be manipulated in an intervention on X. While this issue is a significant one for interventionist accounts of causation, it won't concern the argument in the paper.

argument will bolster the Recognitional Competence Account's non-causal explanation of the crucial counterfactual.

Let's consider some ways in which we might manipulate my knowledge of what Fs look like (call this variable *X*) with respect to my knowledge that those are Fs (call this variable *Y*). Again, we're here assuming that to know what Fs look like is to possess propositional knowledge that Fs are W. Certain ways of manipulating *X* will "switch off" *Y*, not because my knowledge that those are Fs is *based on* my knowledge that Fs are W, but because of metaphysical dependence relations holding among constituents of my knowledge. For example, suppose we manipulate *X* by "switching off" my possession of the concept F. Plausibly, belief is a necessary condition on propositional knowledge, and possessing the concepts that figure in the content of a belief that *p* is necessary for believing that *p*. So switching off my possession of the concept F will be a way of switching off my propositional knowledge that Fs are W. However, possessing the concept F is also a necessary condition for knowing that those are Fs, for the reasons just outlined. If I do not possess the concept F, then I cannot have propositional knowledge that those are Fs.

Suppose now that we follow Woodward's suggestion for how to treat manipulations of variables that bear metaphysical dependence relations to other variables. In that case, we should think of manipulating my knowledge that Fs are W by switching off my possession of the concept F as tantamount to switching off my knowledge that those are Fs. For this reason, the proposed manipulation doesn't count as an intervention. The manipulation automatically switches off *Y*, and so violates the following necessary condition on being an intervention: "I must not directly [change] *Y* via a route that does

not go through X” (Woodward 2016, 20). Thus, the proposed manipulation fails to isolate the causal relevance of my knowing that Fs are W with respect to my knowledge that those are Fs.

We should consider, then, other putative interventions. Helpfully, the supermarket case draws our attention to a range of possible types of interventions by contrasting my epistemic situation with Dmitri’s: If I didn’t know what avocados look like *because I was relevantly like Dmitri*, then I wouldn’t know that those are avocados. Dmitri possesses the concept of an avocado, and responds differentially to experiences as of avocados, but doesn’t know (*de dicto*) that avocados look W. Suppose, then, that we try to intervene on my knowledge that Fs look W (with respect to my knowledge that that’s an F) by switching off my belief that Fs look W.

I’ll argue that this sort of intervention is metaphysically impossible, because of how my belief that Fs are W is metaphysically related to other features of my epistemic life. For ease of explanation, I’ll illustrate the point in terms of credences, but it will apply equally well in the case of full belief, disbelief, and suspension of judgment. Suppose that I initially have a very high credence in L, the proposition that Fs look W:  $Cr(L)=.9$ . But then, suppose, we manipulate my credence to make me less confident in L by giving me some extremely reliable testimony that  $\sim L$ : Fs actually don’t look W. We thereby set my credence to something much lower, say,  $Cr(L)=.1$ .

Before continuing with this line, let me address a worry. Namely, it might seem as though we should try to intervene in a more surgical way on my credence in L than by giving me a piece of evidence. Couldn’t we, for example, more surgically intervene on my credence by changing my brain state? In response, we should note two things. First,

the motivation for the worry may come from a lingering commitment to a Lewisian-style semantics for counterfactuals, according to which we interpret counterfactuals by considering possible worlds most similar to the actual world in which the antecedent is true. On an interventionist picture, however, there is no requirement on interventions that they preserve similarity to the actual world. Second, we want the manipulation on my credence in  $L$  to hold fixed that I'm rational, since we want to use the intervention to assess the causal influence of one piece of knowledge on another for a rational subject. Holding fixed that I'm rational, we cannot intervene on my credence in  $L$  in the right way without changing my evidence (or my priors). Suppose, for example, that we intervene in some way on my brain state, lowering my credence to  $Cr(L)=.1$  without changing my evidence (or priors). In that case, I'm very confident that it's not the case that  $Fs$  look  $W$ , but my total evidence is the same as it is in the supermarket case, and so (given my priors) supports high credence in  $L$ . So to preserve my rationality, we should consider a manipulation that lowers my credence in  $L$  by changing my evidence, e.g., by giving me some reliable testimony that  $\sim L$ .

Now we should ask how setting my credence in this way is metaphysically related to other features of my epistemic life. If we try to intervene by giving me strong enough evidence that  $\sim L$  to lower my rational credence to  $Cr(L)=.1$ , then plausibly we have thereby given me a rebutting defeater for the belief that those are  $Fs$ . After all, I've acquired strong evidence that  $Fs$  don't look  $W$ , and I'm now having an experience (as I sort the  $Fs$ ) as of an object that looks  $W$ . So the evidence I've acquired that  $Fs$  don't look  $W$  constitutes an excellent reason to believe that those aren't  $Fs$ .<sup>12</sup> In virtue of possessing

---

<sup>12</sup> The evidence I acquire may very well have to be stronger than just one piece of reliable testimony to the effect of  $\sim L$ , since that may not be strong enough evidence to constitute a rebutting defeater. We might

a rebutting defeater for the belief that those are Fs, I now no longer count as knowing that those are Fs. Even defenders of immediate perceptual knowledge grounded in recognitional competence can accept that possessing a defeater for the belief that p is sufficient to undermine knowledge that p.

If we accept Woodward's account of how to handle metaphysical dependence in an interventionist framework, we'll want to say that by manipulating my knowledge that Fs look W, and setting my credence to  $Cr(L)=.1$ , we thereby switch off my knowledge that those are Fs. But here again, the manipulation doesn't count as an intervention: It doesn't isolate the causal influence of X on Y, because it directly changes the value of Y. Giving me a rebutting defeater for my belief that Fs are W will cause me to lose my knowledge that Fs are W, but it will also directly cause me to lose my knowledge that those are Fs. Thus, the manipulation violates the following necessary condition on being an intervention: "I must not directly [change] Y via a route that does not go through X" (Woodward 2016, 20).

Of course, there is another way to try to intervene on my knowing what avocados look like to make me relevantly more like Dmitri. We could lower my credence less drastically so that it is set to  $Cr(L)=.5$ , instead of  $Cr(L)=.1$ . Again, we'll want to hold fixed that I'm rational, since we are trying to assess the causal influence of one piece of knowledge on another in a rational subject. So we could try to intervene to lower my credence by changing my evidence (while holding fixed by priors). This time, let's say, I acquire testimony from many reliable testifiers that the manager, who has told me that what I'm sorting are avocados, is grossly inept and misidentifies produce about half the

---

imagine that, in addition to receiving testimony from many reliable testifiers to the effect that  $\sim L$ , I also acquire a compelling explanation of why all of my original evidence for L was misleading.

time. (We're imagining that the manager's testimony was my only evidence for L.)<sup>13</sup> Suppose my rational credence, in light of this evidence, is  $Cr(L)=.5$ . Now I'm not particularly confident or unconfident that Fs look W. What happens to my knowledge that those are avocados under an intervention of this sort?

Plausibly, in this case, I acquire an undermining defeater for the belief that those are avocados. When I learn that the manager may have misinformed me when she testified that the things I'm sorting are avocados, I acquire a defeater that cuts the support relation between my visual experience and the belief that those are avocados. I have a visual experience as of an object that looks a certain way, but now I have reason to believe that the fact that something looks that way doesn't particularly support (or disconfirm, for that matter), the belief that it's an avocado. Again, since a lack of defeaters is a necessary condition on knowledge, even immediate perceptual knowledge, I wouldn't count as knowing in this case that those are avocados. Here too, the manipulation directly switches off my knowledge that those are avocados, and so fails to isolate the causal influence of my knowledge of what avocados look like. Thus, the manipulation isn't an intervention.

## 6. Intervening on Knowledge Via Justification

In the last section, we considered manipulations on X with respect to Y that target my belief or high credence. We saw that the most plausible such manipulations fail to be

---

<sup>13</sup> In this case, my rational credence in L is much higher than Dmitri's. But one might think that the relevant intervention would render my and Dmitri's epistemic situation identical at least as far as L goes. The next section addresses putative interventions like this. Thanks to Matt for the objection.

interventions, and so can't be used to demonstrate a causal dependence between the two pieces of knowledge. In this section, we'll consider manipulations on X with respect to Y that target my justification for the belief that Fs look W. I'll argue that, here too, the most plausible putative interventions fail to be interventions, and so cannot be used to demonstrate causal dependence.

There seems to be a connection between whether my belief that Fs are W is justified, and whether I count as knowing that those are Fs (McGrath 2017, 34). The following case illustrates the dependence:

**Lucky Guess:** Suppose I have no reason to believe that the fruits Dmitri and I are arranging, which have look W, are avocados. I simply guess that they are, and thereby come to believe that avocados look W. My belief is true, but unjustified. I also thereby acquire a competence to recognize avocados: I am now disposed to believe truly that that's an avocado when I have the appropriate sort of experience. The next day, when I see a new shipment of avocados, I form the belief that those are avocados. Let's stipulate that my belief is an exercise of my recognitional competence (Adapted from McGrath 2017, 34).

Two things seem clear about the case. First, I don't *know* what avocados look like; my belief that avocados look W is unjustified and so does not constitute knowledge. Second, when I see the new shipment of avocados, my belief that those are avocados does not constitute knowledge. (We are assuming that, in the meantime, I acquire no extra relevant information.)

Now let's revisit the counterfactual. On the suggestion we are considering, we manipulate the causal history of my belief that Fs look W so that I form the belief by guessing and acquire, by accident, a competence to recognize Fs. (It's not so important here exactly how we ensure that I guess and thereby form a belief; we needn't hold fixed my rationality in this case, as I'm already forming beliefs by guessing.) Is this a *bona fide* intervention on my knowledge of what Fs look like with respect to my knowledge that those are Fs? Again, I think the answer is no. On the proposed manipulation, it is lucky that I acquire a competence to recognize Fs. This fact will be significant on many accounts of immediate perceptual knowledge that appeal to a recognitional competence. I'll review two ways in which the fact that my recognitional competence was acquired in this lucky way explains why my belief that those are Fs fails to constitute knowledge. The first appeals to a necessary condition on the acquisition of methods in process reliabilism, while the second appeals to requirements on knowing full well in virtue reliabilism. The two strategies are compatible. The plausibility of either one will be enough to show why the proposed manipulation fails to constitute a proper intervention. By ensuring that my competence was acquired in this lucky way, the manipulation directly causes me to lose my knowledge that those are Fs, and so doesn't isolate the causal influence of my knowledge that Fs are W.

### *6.1 Method Acquisition*

The lucky guess case is analogous to this case from Goldman (1986):

**Lucky Method Acquisition:** “Suppose Gertrude’s mathematical education is seriously deficient: she has never learned the square root algorithm. One day she runs across the algorithm in a pile of papers written by someone she knows to be a quirky, unreliable thinker, and no authority at all on mathematical matters. Despite this background knowledge, she leaps to the conclusion that this rule for deriving square roots (the rule is so labeled) is a sound method. [Indeed, she’s right: the algorithm is a reliable method for deriving square roots.] She proceeds to follow, and form beliefs in accordance with, this algorithm...Are these beliefs justified? Clearly not...” (91).

These sorts of cases motivate Goldman to introduce the following necessary condition on doxastic justification.

**Method Acquisition Condition:** S’s belief that p, which is the output of reliable method M, is justified only if M is acquired or modified through a metareliable second-order process (1986, 51-53, 91-93).

Second-order processes are processes through which a subject acquires or modifies a belief-forming method. Methods here are learned heuristics, algorithms, or methodologies. Second-order processes will include learning strategies, where what is learned is a new belief-forming method, rather than a proposition. A second-order process is metareliable just when it is sufficiently reliable in outputting reliable first-order

methods (1986, 52-3). No vicious regress threatens because many second-order processes will be learning processes that need not be acquired in the same way as methods.<sup>14</sup>

The method acquisition condition can explain why Gertrude cannot use the algorithm she has learned to form justified beliefs about square routes, even though using the algorithm is a reliable method. The second-order process through which she acquires this method, i.e., trusting someone she knows to be an unreliable, unauthoritative thinker, is not meta-reliable: It is not a reliable way of acquiring reliable methods.<sup>15</sup>

We can make a similar point about our lucky guess case. Even if we leave open the possibility of innate recognitional competences, some recognitional competences will be acquired methods, including plausibly the competence to recognize avocados. Simply guessing about what Fs look like is not a meta-reliable second-order process; it is not a reliable way of acquiring reliable first-order recognitional methods. Even if one gets lucky, and happens to acquire a reliable first-order recognitional method, that method does not yield justified beliefs. Thus, when we manipulate my knowledge of what avocados look like, and ensure that I merely guess that avocados look W, we thereby ensure that the reliable first-order method through which I recognize avocados is not itself acquired in a reliable way. If we accept the method acquisition condition, the

---

<sup>14</sup> Swampman's competence to recognize avocados is unlearned and so plausibly can be understood as a process rather than a method. Thus, I think the method acquisition condition can allow that Swampman can have perceptual knowledge that *that's an avocado*. For the Swampman objection to proper functionalism, see Sosa (1993), and for the Swampman objection to teleological theories of mental content, see Davidson (1987).

<sup>15</sup> Lyons (2009) and Pace (2010) endorse similar views. Lyons (2009) holds that justified basic perceptual beliefs are the outputs of perceptual systems, and that a necessary condition on a perceptual system is that it "has developed as a result of the interplay of learning and innate constraints" (136). Thus, my belief in the Lucky Guess case won't count as a justified basic perceptual belief, since it is not the output of a perceptual system. Likewise, Pace (2010) suggests that only recognitional competences that are "appropriately acquired" can yield justified beliefs (428).

manipulation thus ensures that this first-order method does not yield justified beliefs. My belief that those are avocados, which by stipulation is an exercise of this reliable recognitional method, would then not count as knowledge because it is unjustified.

The proposed manipulation, then, fails to constitute a *bona fide* intervention. The manipulation directly causes me to lose my knowledge that those are avocados, rather than isolating the causal influence of my knowledge that avocados are W.

An objector might point out that the method acquisition condition won't help us handle every possible manipulation of my doxastic justification for L. It seems we can imagine cases in which I acquire a recognitional competence in an appropriate way, but where we "switch off" my doxastic justification for L in some other way, (e.g., perhaps I simply forget L or forget all my evidence for L). Won't it then be intuitive that I will no longer know that those are Fs?<sup>16</sup> My response is to agree that the method acquisition condition won't help with all manipulations on doxastic justification. It just helps with cases like Lucky Guess. However, I suspect that, once the details of the proposed manipulation on doxastic justification are filled in, it will be clear that the manipulation directly impacts on knowledge that those are Fs. For example, if simply forgetting L involves losing belief in L, then considerations from Section 3 come into play. Plausibly, the subject would thereby lose the disposition underlying her competence to recognize Fs. If simply forgetting L involves forgetting any evidence I have for L, (a forgotten evidence case), then it's not as clear that I in fact do lose doxastic justification for L.

## 6.2 *Knowing Full Well*

---

<sup>16</sup> Thanks again to Matt McGrath for pushing this objection.

Even if we reject the method acquisition condition, we have an alternative strategy for showing that the manipulation fails to constitute an intervention. The strategy draws on resources from virtue reliabilism, particularly Sosa's (2015, 2011) account of knowing full well.

Consider, again, what happens when we intervene to ensure that I merely guess what avocados look like. I guess the truth, that avocados look W, and thereby acquire a competence to recognize avocados. Now suppose that I see some avocados and exercise my newfound recognitional competence in judging that they are avocados. On Sosa's view, in *judging* that they are avocados, my aim is more robust than merely judging truly. Rather, a judgment aims at aptness: It succeeds not merely when it is true, but when its truth manifests competence.<sup>17</sup>

**Apt Performance:** S's performance is apt just when it is successful and its success manifests S's competence.

My judgement that those are avocados in the Lucky Guess case counts as successful, then, because it is apt: By stipulation, its truth manifests my ill-gotten competence to recognize avocados.

However, a merely successful judgment can still fall short in an important way and fail to constitute knowing full well. A successful judgment can still fail to attain its aim aptly; that is, the judgment can fail to be *aptly* apt. Consider, for example, a

---

<sup>17</sup> For discussion of full aptness and the aim of judgment, see Sosa (2015, ch.3).

basketball player who is insensitive to the conditions under which her shots are likely to go in; she is just as likely to shoot from across the court as from directly under the basket. When she shoots successfully from directly under the basket, her shot might be apt because its success manifests her first-order competence. However, her choice to shoot there is not guided by sensitivity to the conditions under which her shots would be apt; far too easily would she have made an inapt shot from across the court. Thus, her shot is not aptly apt; its aptness is due to luck rather than guided by second-order competence to shoot in conditions where the shot would be apt (Sosa 2015, 70-73). Likewise, knowledge full well requires not merely that one's judgment be successful, but that it be aptly successful. Since the aim of judgment is apt affirmation, or affirmation whose truth manifests competence, S attains *this aim* aptly only if her affirmation is guided by higher-order sensitivity to the conditions under which the affirmation would be apt. We can capture this point with the following requirement:

**Knowing Full Well:** S's judgment that p constitutes knowing full well only if not too easily would S have judged that p inaptly (Sosa 2015, ch.3).

In the Lucky Guess case, the subject's judgment fails to constitute knowing full well because it violates this condition. Though I happen to guess correctly what avocados look like, I very easily could have guessed falsely that they look X or Y, and so acquired a disposition to respond to experiences as of X-looking objects or Y-looking objects with a false belief that *that's an avocado*. Thus, I very easily could have formed the belief that

*that's an avocado* incompetently, by manifesting one of these unreliable dispositions.<sup>18</sup>

The upshot is that my judgment, though it may manifest my first-order recognitional competence, is not aptly apt. I am not appropriately sensitive to the conditions under which I would judge aptly that those are avocados, since I very easily would have formed that belief incompetently, by exercising a defective disposition. Thus, I do not know full well that those are avocados.<sup>19</sup>

The upshot is that, by manipulating my knowledge that avocados look W in the proposed way, we thereby ensure that it's merely lucky that I acquire a recognitional competence, and that I fail to know full well that those are avocados. Here again, the proposed manipulation fails to constitute an intervention on X with respect to Y. The manipulation directly causes a change in Y, rather than isolating the causal influence of X on Y.

## 7. Conclusions

We've seen that some of the most plausible putative interventions fail to constitute interventions. This result supports the thought that, given metaphysical dependence relations among constituents of the two relevant pieces of knowledge, it is

---

<sup>18</sup> Suppose we imagine that some time has passed since the defective acquisition of the competence and that, in the meantime, I acquire no additional information about how avocados look. Would I still, then, fail to form the relevant beliefs in a way that is insensitive to whether they would be apt? I think so, because without any additional information about the looks of avocados, I would still be forming beliefs in a way that is insensitive to whether they would be competent, (and thus, insensitive to whether they would be apt).

<sup>19</sup> I might still count as possessing animal knowledge, since my true belief does manifest recognitional competence. For those who have a strong intuition that the Lucky Guess subject lacks knowledge of any sort, this might represent a reason to favor the earlier "acquisition condition" strategy.

impossible to intervene on the one with respect to the other; any manipulation that turns off the one piece of knowledge will automatically turn off the other. This argument supports the view that what makes the crucial counterfactual true is not some causal/inferential dependence between the one piece of knowledge and the other. Instead, we can explain the counterfactual by appeal to metaphysical dependence relations holding among constituents of my knowledge of what Fs look like and my immediate knowledge that that's an F, when the latter is explained by my competence to recognize Fs.

Thus, we can defend the possibility of immediate perceptual knowledge against the threat posed by the Only-Because Argument. Immediate perceptual knowledge of a perceived object's kind properties is grounded in recognitional competence.

## Chapter 2

### Seemings are Epistemic Feelings

Appeal to *seemings* in epistemology is widespread. Most famously, phenomenal conservatives hold that if it seems to S that p, then S has *prima facie* justification to believe that p (Huemer 2005, 2007; Tucker 2010; Cullison 2010). Several externalists, too, carve out a special role for seemings that are competently formed or otherwise appropriately related to the facts (Bergmann 2013a; Brogaard 2013b; Sosa 2007, 2015; Pace 2010). However, there is little consensus about what seemings are and why we should be committed to the existence of seemings at all.

In this paper, I develop what I call the epistemic feeling account of seemings. Epistemic feelings are a genus of phenomenally conscious mental states that motivate epistemic performance; familiar examples include the feeling that some information is “on the tip of the tongue,” or about to be remembered, and the eureka experience (or aha! moment). I argue that, in a range of central cases of putative seemings, subjects are undergoing an epistemic feeling called *the feeling of rightness*, (but which I think could also aptly be called a feeling of subjective confidence). Feelings of rightness have been posited to explain when subjects are more or less likely to automatically endorse the output of some Type-1 process in judgement (Thompson, Prowse Turner, & Pennycook 2011; Thompson 2009; Simmons & Nelson 2006). The epistemic feeling account gives a characterization of the content, phenomenal character, and functional role of seemings. Moreover, it clarifies the explanatory work that appeal to seemings enables us to do—namely, by appealing to seemings, we can better explain why subjects in a central range

of cases form the doxastic attitudes they do than we can by appealing to subjects' dispositions to believe.

In the first two sections of the paper, I give some background on the seemings literature and clarify the methodology that I adopt here. Namely, I do not rely on linguistic intuitions about the felicitousness of *seems* expressions because I'm worried about equivocating on *seems*.<sup>20</sup> In Section 3, I introduce the central cases of putative seemings. I develop a characterization of seemings as a genus of mental states by evaluating what the central cases have in common and proposing several necessary conditions on a mental state type *M*'s counting as a species of seeming. After introducing epistemic feelings in Section 4, in Section 5, I give an account of feelings of rightness and argue that they are well-positioned to play the role of seemings in our central cases.

## 1. Two Accounts of Seemings

In this section, I'll distinguish between two families of views of seemings: views according to which seemings are a genus or class of mental state types (compare with emotion), and views according to which a seeming is itself a type of mental state with a distinctive functional role (compare with belief). The question, "What is a seeming?" will have different substance for each of these two families of views.

---

<sup>20</sup> Byerly (2012) argues that certain *seems* expressions are ambiguous. For the famous distinctions among the comparative/noncomparative, phenomenal, and epistemic senses of appearance expressions, see Chisholm (1957) and Jackson (1977).

The standard view of seemings—call it the *Genus Account*—holds that seemings form a genus of mental state types of which there are many species.<sup>21</sup> Thus, there needn't be a functional role that all species of seemings share. The question, "What is a seeming?" has two important readings for proponents of the Genus Account. First, it could be asking for the common features that are shared among the species of seeming. These could be either necessary conditions on being a seeming or properties that contingently tend to cluster together in seemings. On a second reading, the question asks whether *each of the species* of seeming is reducible to another type of mental state. For example, suppose we say that philosophical intuition is a species of seeming. In that case, to ask what a seeming is could be to ask whether philosophical intuition is a *sui generis* propositional attitude, a mere defeasible disposition to believe, or neither of these. Moreover, there is no antecedent reason to expect that each of the species of seeming must be treated in the same way; it could be that one species of seeming is reducible to a defeasible disposition to believe, while another is not.

According to a second family of views—call it the *Type Account*—a seeming is not a genus of mental state types but is itself a mental state type.<sup>22</sup> A seeming, on this picture, has its own distinctive functional role *qua* seeming. The question, "What is a seeming?" for the Type Account has a different force than it did above. In this context, the question asks for a characterization of this functional role, and relatedly, whether seemings are best understood as a *sui generis* propositional attitude, a mere disposition to believe, or

---

<sup>21</sup> See, for example, Huemer (2001, 2005, 2007), who is explicit about treating seemings as a genus, e.g., (2007, 30). Tolhurst (1998) and Tucker (2010) are less explicit on this point but are probably most charitably interpreted as thinking of seemings as a genus of mental state types.

<sup>22</sup> Cullison (2010), Skene (2013), and Bergmann (2013) certainly sometimes talk as though they endorse this sort of view, but I leave open whether they, too, are most charitably interpreted as endorsing the Genus Account.

something else. An answer to this question may not have any implications for the nature of another mental state type—e.g., philosophical intuition, apparent memory, etc.—since seemings needn’t be identified with any of these and aren’t here understood as a genus including these.

The distinction between the Genus and Type Accounts of seemings is important. Some commonly-made points in the literature on seemings invite confusion because they do not clearly make this distinction. For example, it’s sometimes said that the standard view of seemings is that they are *sui generis* propositional attitudes.<sup>23</sup> However, this could be read in either of the following ways: first, that seemings are a genus of mental state whose species—intuition, apparent memory, etc.—are all *sui generis* propositional attitudes, or second, that there is a distinct mental state type called a seeming that is itself a *sui generis* propositional attitude. These proposals are importantly different. For example, a challenge facing the Genus Account but not the Type Account is to show that grouping together the species of seemings under a single genus is non-arbitrary and informative. And a challenge facing the Type Account but not *necessarily* the Genus Account is to argue that we should recognize a new type of mental state with a distinctive functional role.

## 2. The Ontological and Semantic Theses

---

<sup>23</sup> Werner, e.g., renders the standard view as, “Seemings are a *sui generis* propositional attitude toward some content p” (2014, 1762).

It will also be helpful to clarify how we will be treating *seems* expressions in the paper, e.g., “It seems to me that p” and “That seems to be G.” To start, let’s distinguish between the following two theses, following Hawthorne and Lasonen-Aarnio (ms):

**Semantic Thesis:** “At least some ordinary uses of the ‘it seems that’-construction are used to report that” a subject is tokening a seeming (4).

**Ontological Thesis:** “There are seemings, a class of *sui generis* mental states with a distinct phenomenology” (4).

Here the ontological thesis expresses a version of the Genus Account of seemings—seemings are a “class” of mental states. Moreover, it expresses a version of the Genus Account according to which the species of seeming are not reducible to other mental states but are *sui generis* propositional attitudes. Let’s interpret the ontological thesis as being neutral on whether seemings are a natural kind. (Whether seemings are a natural kind will depend on how demanding the conditions for natural kindhood in general are, which will determine whether any psychological kinds are natural.)

The semantic thesis holds that some ordinary uses of *seems* expressions report that a subject is tokening a seeming.<sup>24</sup> For example, when I seem to remember that you were at the party, but know that it is a false memory, I can say, “I seem to remember that you were there, but I know that you really weren’t.” According to the semantic thesis, I thereby report tokening a seeming that you were at the party. The semantic thesis allows

---

<sup>24</sup> For discussion of the semantic thesis, see Brogaard (2013a), Byerly (2012), and Cullison (2013).

that some ordinary uses of *seems* expressions are not seemings reports. When I say, for example, “It seems to me that Sanders could defeat Trump in the general election,” I needn’t be expressing a seeming state. In terms of Chisholm (1957) and Jackson’s (1977) distinctions, the latter case is an epistemic use of “seems” in which I’m reporting a belief, or that the proposition is sufficiently likely on my evidence.

The truth of the semantic thesis would lend some support to the ontological thesis. However, I do not think that the ontological thesis stands or falls with the semantic thesis. First, neither thesis entails the other. The semantic thesis doesn’t entail the ontological thesis because it is possible that the relevant uses of *seems* expressions fail to refer to anything. The ontological thesis doesn’t entail the semantic thesis because appeal to seemings could be a useful technical concept, underwriting informative explanations in epistemology, even if *seeming* is not a folk psychological concept, and there is no ordinary use of seems expressions that makes a seemings report. Second, there are ways to argue for the ontological thesis that do not depend on linguistic intuitions about *seems* expressions. For example, one could sketch an account of seemings by identifying the important similarities shared among its species and point to some informative generalizations that appeal to seemings enables us to make.

In fact, for our purposes here, I’ll endorse Hawthorne and Lasonen-Aarnio’s view about the semantics of seems expressions. On their view, the expression “it seems to S that p” is “used to state that the contextually relevant subject (or subjects) thinks that p is sufficiently likely on a contextually relevant body of his or her evidence (where what counts as sufficient is set by context)” (ms, 8). The contextually relevant body of evidence could sometimes include just a token seeming. However, it often will include

much more than that. And even when the contextually relevant body of evidence includes just a token seeming, “it seems to S that p” is best understood as reporting that p is sufficiently likely for S, given that body of evidence, rather than simply that the subject is tokening the seeming. Thus, on this view, the expression is not best understood as a seemings-report. In the argument that follows, then, I won’t take the felicitousness of the expression, “it seems to S that p” to indicate that S is tokening a seeming.

Hawthorne and Lasonen-Aarnio suggest that the semantic and ontological theses depend closely on each other. They argue that the falsity of the semantic thesis is good evidence for the falsity of the ontological thesis, writing:

If the ontological thesis were true, shouldn’t we expect there to be a word in natural language which picks out seeming-states, just as we have words like ‘pain’, ‘fear’, and ‘belief’? Perhaps we shouldn’t expect there to be natural language words for every distinct kind of mental state, but seeming-states are supposed to play a crucial role in our epistemic lives. (ms, 10)

It seems that this point is targeted against a Type Account of seemings, according to which seemings are a type of mental state with a distinctive functional role. After all, Hawthorne and Lasonen-Aarnio here compare seemings with mental states like pain, fear, and belief. However, I think the point is less effective against a Genus Account of seemings. It might be that the concept *seeming*, as a genus of mental state types, is a technical concept that is useful for making informative generalizations in epistemology. As such, we shouldn’t necessarily expect there to be a folk psychological concept *seeming*, or a natural language expression that reports that one is tokening a mental state that belongs to this genus. (Of course, this is all compatible with saying that there are folk psychological concepts of the various *species* of seeming and natural language

expressions that can be used to report that one is tokening one of the species of seeming.) Thus, the falsity of the semantic thesis doesn't tell strongly against the ontological thesis, so long as the latter is best understood as expressing a version of the Genus Account of seemings.

### 3. What are Seemings?

In this section, I'll present a characterization of seemings without relying on linguistic intuitions about the felicitousness of *seems* expressions. I'll introduce four cases and argue that the subjects' mental states in each case share important features such that they are fruitfully understood as members of a genus. I'll consider a natural alternative explanation of the cases that doesn't appeal to seemings and argue that it fails to capture an important common feature of the cases, which the seemings account can capture.

I'll start by introducing the cases using neutral language:

**Apparent Memory:** You believe that the capital of Missouri is St. Louis. But then you read on the state of Missouri's website that this is a common misconception: The capital is actually Jefferson City. Now, when you try to remember the capital of Missouri, both cities come to mind, but you judge correctly that the capital is Jefferson City.<sup>25</sup>

---

<sup>25</sup> Adapted from Weisberg (forthcoming, 28).

**Cognitive Reflection Test:** You are taking the cognitive reflection test and try to answer the following question: “If it takes 5 machines 5 minutes to make 5 widgets, how long does it take for 100 machines to make 100 widgets.” The answer 100 minutes immediately comes to mind, and you select that answer.<sup>26</sup>

**Competing Intuitions:** Upon reading the Swampman case for the first time, you are torn in two directions. When you focus on the similarities between Swampman and ordinary subjects, you are tempted to say that his beliefs can be warranted. But when you focus on the fact that Swampman’s mental systems lack proper function, you are tempted to deny that his beliefs can be warranted. Both intuitions are equally strong, and so you suspend judgment on the issue.

**Mineral Identification:** You are a geology student identifying samples of minerals for an exam. The first mineral is translucent and its edges break in a regular pattern. Calcite and gypsum both come to mind, but you are unsure of which it is and suspend judgment on the issue.

The proposal is that something similar is happening in each of these cases: Each subject tokens one or more *prima facie* seeming states, and these seemings help explain which doxastic attitudes the subjects form.<sup>27</sup> Consider, for example, the final two cases,

---

<sup>26</sup> The question appears in Frederick (2005). For discussion of the role “cognitive illusions” play in dual-process theories of cognition, see Kahneman and Frederick (2002).

<sup>27</sup> On the distinction between *prima facie* and *ultima facie* seemings, see Tolhurst (1998). He suggests that the distinction is analogous to Ross’s (1930) distinction between *prima facie* and *ultima facie* obligations: A subject might be *prima facie* obligated to  $\Phi$  and *prima facie* obligated to  $\Psi$ , though it is impossible for the subject both to  $\Phi$  and  $\Psi$ . What the subject has an *ultima facie*

which are similar. In Competing Intuitions, the subject tokens a *prima facie* seeming that Swampman can have warranted beliefs and a *prima facie* seeming that he cannot. The subject suspends judgment because the seemings are equally strong or compelling, and (let's suppose) she has no other relevant information. Likewise, in Mineral Identification, it *prima facie* seems to the subject that the mineral is calcite, and it also *prima facie* seems to her that the mineral is gypsum. Again, she suspends judgment because the seemings are equally strong or compelling, and she has no other relevant information.

Likewise, in the first two cases, the subjects' doxastic attitudes are partly explicable in terms of the seemings they token. In Cognitive Reflection Test, it *prima facie* seems to the subject that the answer is 100 minutes. There is no competing seeming in the case, and the subject doesn't engage in the conscious reasoning that might produce the correct answer. Thus, the subject judges that the answer is 100 minutes. Finally, in Apparent Memory, the subject tokens conflicting *prima facie* seemings: that the capital is Jefferson City and that the capital is St. Louis. But because the subject knows about the common misconception, she discounts the latter seeming and judges that the capital is Jefferson City.

Why should we think that the subjects' mental states in each case are species of one genus? I'll here highlight a number of features that they have in common before explaining each one. I'll present these features as necessary conditions for mental state type M to count as a *prima facie* seeming.

---

obligation to do depends on the strength of the competing *prima facie* obligations. Likewise, how things seem to the subject all-things-considered depends on how things seem to her *prima facie*. Sosa (2015) also appeals to this distinction. I agree with Tolhurst that *prima facie* seemings are defeasible, and that it can be rational for a subject to token *prima facie* seemings with incompatible contents. But I'm officially neutral here on how to understand all-things-considered seemings; the epistemic feelings account just concerns *prima facie* seemings.

**Content:** A mental state type M is a *prima facie* seeming only if M has propositional content.

**Direction of Fit:** A mental state type M is a *prima facie* seeming only if M has a mind-to-world direction of fit.

**Defeasibility:** A mental state type M is a *prima facie* seeming only if M is defeasible.

**Phenomenal Character:** A mental state type M is a *prima facie* seeming only if M has presentational phenomenal character.

**Incompatible Content:** It can be rational for a subject to token *prima facie* seemings with incompatible contents at the same time, or more carefully:

A mental state type M is a *prima facie* seeming only if there is some token of M, m, with content p, and some distinct token *prima facie* seeming, s, with content q, where p and q are incompatible, such that it is rational for S to token m and s at time t.

The content condition holds that a mental state type M is a *prima facie* seeming only if M has propositional content. What seems true in each case is a proposition, e.g., that the

capital of Missouri is Jefferson City, that Swampman can have warranted beliefs, that the answer is 100 minutes, and that the mineral is calcite. The content condition officially leaves open whether the propositional content of seemings is conceptually structured, i.e., whether seemings have same general type of content as thought or judgment. However, as I'm imagining these four cases, the subjects' seemings do have conceptually structured content that could be directly endorsed in a judgment.

Like belief, *prima facie* seemings have a mind-to-world direction of fit. They succeed when they accurately represent the world. My seeming that the capital of Missouri is St. Louis is defective because it's inaccurate, unlike a desire, which is not defective just because it is unfulfilled.

Moreover, a mental state type M is a *prima facie* seeming only if M is defeasible. The only case above in which a subject's *prima facie* seeming is actually defeated is Apparent Memory: S tokens a *prima facie* seeming that St. Louis is the capital of Missouri despite knowing that this is incorrect and believing that Jefferson City is the capital. However, the subject in each of the cases could acquire a defeater for any of their *prima facie* seemings.

According to the phenomenal character condition, a mental state M is a *prima facie* seeming only if M has presentational phenomenal character. This condition will take some unpacking. The notion of presentational phenomenal character is often introduced through contrasts. Desiring that you rescheduled the flight for 10am rather than 7am, or imagining that you have done so, has a different phenomenal character than suddenly remembering that you rescheduled the flight. The latter feels as though it is presenting its content as being true, while the former states of desire and imagination do

not. This presentational phenomenal character is variously described as “forceful” (Huemer 2011, 77-79) or “assertive” (Tucker 2010, 53). Tolhurst calls it “felt veridicality,” or the feel of a state that purports to reveal how the world is (1998, 299).

Sometimes the presentational phenomenal character of seemings is described metaphorically as “pulling” subjects towards forming a belief. For example, Sosa (2007) writes:

What are these seemings? It is helpful to compare deliberation on a choice or the pondering of a question, where we “weigh” reasons pro or con. Switching metaphors, we feel the “pull” of conflicting considerations. No matter the metaphor, the phenomenon is familiar to us all. There is something it is like to feel the pull of contrary attractions as we deliberate or ponder. (47)

It is important that we dispel a possible misinterpretation of this metaphor. The relevant phenomenal character shouldn’t be identified with that of a felt inclination to believe. It is a familiar point that, since *prima facie* seemings are defeasible, it is possible to token a *prima facie* seeming without feeling at all inclined to believe its content. We can imagine the Apparent Memory case in that way. I might not feel at all inclined to believe that the capital of Missouri is St. Louis, once I learn that this is a common misconception, even though I token a seeming with that content when that answer comes to mind.

Nevertheless, the seeming has presentational phenomenal character: Unlike desire and imagination, it has the feel of a state that presents its content as being true or accurate. Thus, the presentational phenomenology of a seeming shouldn’t be identified with the phenomenal character of a felt inclination to believe.

According to the incompatible content condition, it can be rational for a subject to token *prima facie* seemings,  $s_1$  and  $s_2$ , with incompatible content at the same time. This is what is happening in Apparent Memory, Competing Intuitions, and Mineral

Identification. For example, the student tokens a *prima facie* seeming that the mineral is calcite and a *prima facie* seeming that the mineral is gypsum. Moreover, it is plausible that she is not irrational or otherwise negatively epistemically evaluable because she tokens both seemings. (The incompatible content condition does not say that it can be rational for a subject to token a *prima facie* seeming directed towards a contradiction.)

Thus, the mental states of the subjects in each case share important features: They are defeasible mental states with propositional content, a mind-to-world direction of fit, and presentational phenomenal character. Moreover, they're governed by the norm given in Incompatible Content.

At this point, it will be helpful to consider two objections. First, one might worry that the necessary conditions on seemings don't do much to delineate the genus; after all, they are just necessary conditions. In particular, credences might seem to satisfy the necessary conditions. Credences have propositional content, and perhaps they can be understood as having mind-to world direction of fit. It can be rational for a subject to have some credence in incompatible propositions, so long as she obeys the axioms of the probability calculus. One might distinguish, moreover, between dispositional and occurrent credences, and hold that occurrent credences have a presentational phenomenal character that corresponds to how confident someone feels in some proposition.

In response, I think it is more difficult for credences to satisfy the necessary conditions on seemings than it seems at first. Even if one distinguishes between dispositional and occurrent credences, it is an open question whether the best way to make the distinction is to hold that occurrent credences have phenomenal character. An alternative would be to say that occurrent credences, like occurrent full beliefs, are poised

to play a role in guiding action. It is also difficult to understand what it would mean for credences to be defeasible in the same way that seemings are. It can be rational for a subject to token a *prima facie* seeming that *p*, even when she knows that *p* is false. But if a rational subject knows that *p* is false, then plausibly  $Cr(p)=0$ . Credences are not particularly apt, then, to play the role of seemings.

There is a more pressing issue lurking here. We should countenance seemings only if we must appeal to them in the best explanation of some psychological or epistemological phenomenon. But it might seem as though there are plausible alternative explanations of our central cases that do not appeal to seemings. A natural alternative explanation proceeds in terms of dispositions to believe. In the Cognitive Reflection Test, the subject doesn't token a *prima facie* seeming that the answer is 100 minutes; she is simply disposed to form that belief (in the absence of defeaters). Likewise, the subject in the Apparent Memory case doesn't token a *prima facie* seeming that the capital of Missouri is St. Louis. Rather, she is disposed to believe that proposition in the absence of defeaters, but because she possesses a defeater, she does not form the belief. Since we are already committed to the existence of dispositions to believe, we should explain the cases in those terms. Moreover, we can't fall back on linguistic intuitions about *seems* expressions here.

Before presenting a response to the objection, I want to highlight one difficulty. It is somewhat unclear exactly how to explain the Competing Intuitions and Mineral Identification cases on a dispositionalist account. How exactly are we to understand the subjects' dispositions to believe in these cases? The subjects are not disposed to believe contradictions, e.g., that the mineral is both calcite and gypsum, or that Swampman both

can and cannot have warranted beliefs. Nor is it the case that the subjects possess defeaters, as in the Apparent Memory case. So we cannot say simply that the subject has a disposition to believe, e.g., that the mineral is calcite that is triggered by the absence of defeaters. To handle these cases on a dispositionalist picture, we have to say something complicated about the trigger condition for the relevant dispositions to believe. (Again, this is to account for the fact that the subject would only believe that, e.g., the mineral is calcite if it didn't also *prima facie* seem to her that the mineral is gypsum.) Let's suppose, however, that we can work out some way of expressing the relevant dispositions for the latter two cases.

There is still a problem with the dispositionalist account. It is certainly true that our subjects have complex dispositions to believe, but these dispositions don't exhaust what is interesting about these cases: that subjects are in states that, as a matter of their phenomenal character, seem to present or reveal some proposition as being true. Dispositions to believe don't seem apt for explaining this feature of the cases. It is false that S has a disposition to believe p just in case S tokens some mental state with presentational phenomenal character with respect to p. I can have a disposition to believe (or judge) that giraffes don't wear pajamas even if I've never considered the proposition before and am not tokening a mental state with presentational phenomenal character with respect to that proposition.<sup>28</sup> The dispositionalist account of the central cases is at least incomplete.

---

<sup>28</sup> Werner (2014) defends the view that a seeming that p is *token* identical with a disposition to believe p. The view is compatible with holding that S can possess a disposition to believe p without undergoing presentational phenomenal character with respect to p. However, because the view gives a token identity, it doesn't obviate the need to give an account of seemings as a mental state type or genus.

I suspect that some of the resistance to characterizing our central cases in terms of seemings owes the fact that seemings are still fairly mysterious. In what follows, I'll try to make seemings less mysterious by developing the epistemic feelings account and applying it to our central cases. I'll argue that our subjects can be understood as tokening feelings of rightness, a species of epistemic feeling.

#### 4. Epistemic Feelings

Since epistemic feelings are not widely discussed in philosophy, I'll briefly introduce them here before characterizing feelings of rightness. Consider the following examples of epistemic feelings:

**The Tip of the Tongue Feeling:** You are struggling to remember the name of the author of *The Glass Menagerie*. You are confident that his name starts with a "T," that he was an American playwright, and that he also wrote *A Streetcar Named Desire*, but you can't quite produce his name. It feels as though the name is on the tip of your tongue.

**The Feeling of Knowing:** You are a contestant on a gameshow in which you have to push a buzzer faster than the other contestants to answer the question. You will lose points if you push the buzzer but get the answer wrong. The host asks for the capital of Honduras. You feel as though you know the answer, and rush to press the buzzer. It takes you several seconds before you can recall the answer.

**The Feeling of Understanding (Eureka Experience or Aha! Moment):** You're an Intro to Logic student who is struggling to understand why arguments that affirm the consequent are invalid. Your instructor gives you a particularly helpful example of such an argument. After considering the new example, suddenly everything "clicks" and you feel as though you understand why such arguments are fallacious.

**The Feeling of Forgetting:** Leaving a room, you're overcome by the feeling that you have forgotten to do something important that you had intended to do in that room.<sup>29</sup>

Epistemic feelings are a genus of mental states with three central properties. First, they are phenomenally conscious. Second, part of their functional role is that they tend to motivate epistemic performances. Third, it is plausible that they are sensitive to subpersonal monitoring of other mental processes. I'll address each of these in turn.

Epistemic feelings have phenomenal character. There is something it's like to have the answer on the tip of one's tongue, to feel that one knows the answer, to suddenly understand something in an aha! moment, or to feel as though one has forgotten something. We might wonder, then, whether epistemic feelings have cognitive phenomenology, or phenomenal character that is neither sensory nor affective. It is clear that the epistemic feelings considered here don't have sensory phenomenal character. I

---

<sup>29</sup> For these and additional examples of epistemic feelings in the philosophical literature, see Arango-Muñoz (2014), Michaelian and Arango-Muñoz (2014), Proust (2008, 2013), and Dokic (2012, 2014).

think it is an open question, however, whether the phenomenal character of epistemic feelings is best understood as affective. On one view, epistemic feelings are reducible to ordinary emotional feelings that are accompanied by (or whose content concerns) cognitive events. The eureka experience, for example, might be an ordinary feeling of excitement that simply accompanies the event of acquiring understanding or knowledge, while the tip of the tongue feeling and feeling of knowing would be ordinary feelings of anticipation directed at remembering something. Likewise, perhaps the feeling of forgetting is an ordinary feeling of uneasiness that accompanies the thought that one has forgotten something. My focus here won't be on evaluating these proposals; the important point is that recognizing epistemic feelings doesn't require a commitment to cognitive phenomenology.<sup>30</sup>

To see the sense in which epistemic feelings are motivational, let's revisit the examples. In the first, your tip of the tongue feeling motivates you to continue trying to remember the name of the author: You feel as though you are about to remember it and that if you keep trying, you'll succeed. In the second, your feeling of knowing the answer motivates you, first, to press the buzzer, (an ordinary intentional action), and, second, to try to recall the capital of Honduras. In the third, your feeling of understanding motivates you to end inquiry, (in this case, on the question of why affirming the consequent is a fallacy). Finally, your feeling of forgetting might motivate you to try to remember what you were planning to do. Because they are motivational, epistemic feelings play a role in regulating epistemic activity. Epistemic feelings can motivate initiating a performance,

---

<sup>30</sup> See Arango-Muñoz (forthcoming) for an argument that the phenomenal character of the tip-of-the-tongue feeling in particular is not a case of cognitive phenomenology. For discussion of the controversy over cognitive phenomenology in general, see the essays in Bayne and Montague (2011).

(e.g., initiating a retrieval attempt), continuing with some performance, (e.g., continuing the retrieval attempt), or ending some performance, (e.g., ending inquiry on some question) (Arango-Muñoz 2014).

Finally, it is plausible that the mnemonic epistemic feelings—the feeling of knowing, feeling of forgetting, and tip of the tongue feeling—are responsive to subpersonal metacognitive monitoring processes. These monitoring processes are metacognitive in the sense that they are sensitive to the properties of other mental events and processes. For example, both the tip of the tongue feeling and the feeling of knowing are sensitive to properties of your retrieval attempt. Both feelings are stronger when the conditions that elicit the retrieval attempt are familiar, (e.g., when you remember having seen a certain question before), and when you can recall ancillary pieces of information about the target item, (e.g., what letter it starts with, or what other words it sounds like).<sup>31</sup> These monitoring processes are automatic, involuntary, and effortless. Though they produce a phenomenally conscious epistemic feeling, the processes themselves are not accessible to introspection.

One might wonder whether there is another sense in which epistemic feelings are metacognitive, namely, whether they have metacognitive content, or content that refers to mental states, events, and processes. Whether epistemic feelings have metacognitive content is a choice point for an account of epistemic feelings. Importantly, epistemic

---

<sup>31</sup> On the causal basis of the feeling of knowing and tip of the tongue feeling—and the related “feeling of familiarity”—see Koriat (1993, 1994, 2000), Reder (1987, 1988, 1996), Reder and Ritter (1992), Schwartz (2002), Schwartz and Metcalfe (2002), Whittlesea (1993), and Whittlesea and Williams (2001). There is also a vast literature on the causal basis of the eureka experience or aha! moment, in which metacognitive monitoring is *not* as central; see, e.g., MacGregor et al. (2001), Knoblich et al. (1999), Bowden et al. (2005). Sprugnoli et al. (2017) gives a helpful review.

feelings could be responsive to metacognitive monitoring even if they do not have content that refers to other mental states, events, and processes.

I'll here review two alternatives for how to understand the content of epistemic feelings. On the first option, all epistemic feelings have conceptually-structured metacognitive content. For example, the feeling of understanding, on this view, represents that *I understand that p*; the feeling of knowing might represent something like *I can remember the answer*; the feeling of forgetting might represent that *I've forgotten something*. To token an epistemic feeling, on this view, requires the possession of concepts of mental states and processes such as *believe*, *remember*, and *understand*. Animals and young children may then lack the conceptual resources necessary to token epistemic feelings.

On another possibility, which is the one that I favor, we can allow that some epistemic feelings have metacognitive content, while others do not. This view is motivated by the thought that epistemic feelings are a genus of diverse mental states, and that we do not have antecedent reason to think that the content of each species must be treated in the same way.<sup>32</sup> We might understand the eureka feeling, for example, as representing, not that I understand that p, but simply that p. Moreover, in the next section, I present an account of feelings of rightness, which are most plausibly directed at first-order propositional content. However, we needn't hold that all epistemic feelings have first-order content. It is less clear what the first-order alternative would be for the tip of the tongue feeling and the feelings of knowing and forgetting. Since these are all cases in which the subject has not yet retrieved the relevant piece of information, it seems

---

<sup>32</sup> Thanks to Frankie Egan for this point.

implausible that that piece of information could be the content of the feeling. While there is some putative evidence that non-human animals, who do not possess sophisticated mental state concepts, token epistemic feelings, it concerns “feelings of uncertainty” rather than these meta-memory experiences.<sup>33</sup> Thus, we can leave open the possibility that these meta-memory experiences have second-order content.

## 5. Feelings of Rightness

### *5.1 Introducing Feelings of Rightness*

In recent work on dual process theories of cognition, some psychologists have posited an epistemic feeling called the “feeling of rightness” (FOR) to explain when subjects are likely to engage in Type-2 reasoning.<sup>34</sup> In this section, I’ll describe this feeling—its functional role, content, and phenomenal character—before arguing that it is well-positioned to play the role of seemings in our central cases.

In tokening a feeling of rightness, a subject feels to a greater or lesser degree that some content is true. We can think of the feeling of rightness as a feeling of subjective confidence, but one that is produced automatically rather than through deliberation on one’s evidence. On Thompson et al.’s (2011) account, the feeling of rightness is metacognitive because it is determined in part by subpersonal monitoring of the

---

<sup>33</sup> For an argument that non-human animals can token feelings of uncertainty, see Proust (2013, ch. 5) and for a challenge to the argument that I take to be compelling, see Carruthers (2017).

<sup>34</sup> See, e.g., Thompson et al. (2011), Thompson (2009), Simmons & Nelson (2006). My focus will be on arguing that FORs are well-suited to play the role of seemings, rather than evaluating the evidence offered in these three papers for positing FORs.

“fluency” of other processing. Some processing is more fluent to the extent that it quickly produces some conscious output, it doesn’t produce competing answers, and more ancillary information comes to mind. While fluency “monitoring” might suggest some conscious, deliberate process, as before the thought is that monitoring processes that yield FORs are subpersonal, automatic, and effortless.

FORs play an important role in certain dual process theories of cognition. On Thompson et al.’s (2011) account, a feeling of rightness is one of the determinants of whether a subject engages in Type-2 reasoning or simply endorses the output of some Type-1 process in judgment. The distinction between Type-1 and Type-2 processes is fraught.<sup>35</sup> For our purposes, we’ll say that Type-1 processes are fast, unconscious (not accessible to introspection, although their outputs might be), automatic, and effortless. Type 2 processes are slow, conscious (accessible to introspection), deliberate, effortful, and operate serially because they require access to working memory.

The proposal is that a relatively fluent Type-1 process produces a strong FOR. All else equal, the stronger the FOR, the more likely it is that the subject will endorse the content of the FOR in judgment without engaging in Type-2 reasoning. A relatively disfluent Type-1 process produces a weak FOR, making it more likely that the subject will engage in Type-2 reasoning. Engaging in Type-2 reasoning doesn’t guarantee that the subject will end up rejecting the content of the FOR in judgment, but it makes that outcome more likely. Moreover, the FOR isn’t the sole determiner of whether the subject engages in Type-2 reasoning on this view; other factors are relevant, including the time allotted to a task and “global characteristics of the reasoner” (Thompson et al. 2011, 108).

---

<sup>35</sup> For an overview of some of the controversies, see Evans (2009).

Suppose, for example, that a subject is playing a trivia game in which she is trying to remember the capital of Iceland.<sup>36</sup> Reykjavík immediately comes to mind and isn't accompanied by a competing answer. Moreover, the subject can remember lots of ancillary information about Reykjavík: that it is on the sea, is the most northerly capital in the world, etc. The proposal is that the subject experiences a strong FOR, (or feeling of subjective confidence), directed towards the proposition that Reykjavík is the capital, and this makes it more likely that she will automatically endorse that content in judgment. Suppose, instead, that both Reykjavík and Helsinki come to mind after a great deal of effort, accompanied by no ancillary information. The subject might then experience a weak FOR that Reykjavík is the capital and a weak FOR that Helsinki is. The weakness of the FOR, on this view, makes it more likely that the subject will engage in Type-2 reasoning. She might, for example, reason consciously about which word is more Icelandic-sounding, before giving up and suspending judgment.

On this picture, the content of the FOR is determined by the Type-1 process whose fluency is monitored. Since at least some of the Type-1 processes that are monitored in this way involve propositional content, (e.g., retrieving a semantic memory), at least some FORs will be directed at propositional content.<sup>37</sup> In the case

---

<sup>36</sup> For this case, see Weisberg (forthcoming), who appeals to epistemic feelings involved in meta-memory to give an account of the relationship between credence and full belief. For the reasons spelled out in Section 3, however, I think we should distinguish credences from *prima facie* seemings and feelings of rightness.

<sup>37</sup> I'll remain neutral here on whether FORs can ever have non-propositional content. In that case, the FOR would be a subjective feeling of confidence that some content is accurate, rather than true. Since plausibly the content of the FOR depends on the type of first-order process whose fluency is monitored, whether the FOR can have non-propositional content depends on which processes are monitored in this way. There are a wide range of effects in which the fluency of mental processing influences the judgments subjects form. For reviews, see Alter and Oppenheimer (2009) and Unkelbach and Greifeneder (2013). However, which of these effects is

above, for example, the content of the FOR is *that the capital is Reykjavík*. Thus, though the processes that yield FORs are sensitive to the properties of first-order processes, we needn't think of them as having second-order content, or content that refers to other mental states or events as such.

As a feeling, the FOR is a phenomenally conscious mental state. Like other phenomenally conscious mental states, the FOR may often go unnoticed or unaccessed. Moreover, the phenomenal character of the FOR is presentational rather than cognitive or imaginative: It has the feeling of a state that reveals how things are. The phenomenal character of the FOR also has intensity: One can feel more or less subjectively confident in some proposition. Importantly, the intensity of the feeling of rightness is not determined by the subject's total evidence or the reliability of first-order processing. Rather, the intensity is determined by the fluency of the first-order process that is monitored. Thus, the intensity of the FOR needn't track any sort of epistemic value, since unreliable processes, and processes that are not sensitive to the subject's total evidence, can be highly fluent.<sup>38</sup> FORs are tied to particular instantiations of first-order processes whose fluency is monitored.

## 5.2 *Feelings of Rightness and Seemings*

---

best explained by supposing that subjects' judgments are shaped by metacognitive *experiences* is a subtle empirical issue that I won't take up here.

<sup>38</sup> In fact, fluency effects are implicated in a number of notorious cognitive biases. See, e.g., Begg, Anas, and Farinacci (1992), Briñol, Petty, and Tormala (2006), Kelley and Lindsay (1993), and Reber and Schwartz (1999). The epistemic feeling view thus raises the familiar worry that seemings can be formed in epistemically defective ways, as in certain cognitive penetration cases. See, e.g., Markie (2005, 2006) and Siegel (2012).

Feelings of rightness satisfy all of the necessary conditions on seemings identified in Section 3. Moreover, as I'll argue, they are apt for explaining the role that seemings play in determining subjects' doxastic attitudes in our central cases.

Feelings of rightness plausibly meet the content, phenomenal character, direction of fit, and defeasibility conditions on seemings. FORs paradigmatically have propositional content. Moreover, they have presentational phenomenal character: They have the feel of states that reveal the way the world is. Given the role they play in shaping judgment—the stronger the FOR, the more likely it is that a subject will simply endorse its content in judgement—it is plausible that they have a mind-to-world direction of fit, like a belief. Finally, FORs are defeasible. It's not the case that tokening an FOR that *p* is sufficient for it's being rational for a subject to believe *p*; *S* could possess a defeater that, e.g., her FORs are generally unreliable.

FORs also satisfy the incompatible content condition. Again, that condition says roughly that it can be rational for a subject to token *prima facie* seemings with incompatible contents at the same time. First, setting aside what is rational for a moment, it is possible for a subject to token FORs with incompatible contents. In trying to remember the capital of Iceland, I might token an FOR that the answer is Reykjavík and an FOR that the answer is Helsinki. Moreover, I don't seem irrational in virtue of tokening both FORs. Thus, the FOR satisfies each of the necessary conditions on seemings sketched in Section 3.

The subjects our central cases can be fruitfully understood as tokening FORs. Consider, for example, the Cognitive Reflection Test case, in which it *prima facie* seems to the subject that the answer is 100 minutes. The thought is that some Type-1 process

fluently produces this answer, yielding a strong FOR directed towards that content. The subject is then likely to endorse that content in judgment without engaging in Type-2 reasoning. Thus, the subject of the case judges that the answer 100 minutes.

In the Mineral Identification case, the subject tokens a *prima facie* seeming that the mineral is calcite and a *prima facie* seeming that the mineral is gypsum. The proposal is that the *prima facie* seemings are FORs that are produced as byproducts of Type-1 processing involved in trying to recognize the mineral. On this picture, FORs are weaker when the Type-1 processing that is being monitored yields competing answers, as in this case. And the weaker the FOR, the less likely it is the subject will automatically endorse its content in judgment without engaging in Type-2 reasoning. Thus, our subject doesn't automatically endorse the content of either FOR, but rather suspends judgement. We can tell a similar story about the Competing Intuitions case, which is structurally similar.

Finally, in the Apparent Memory case, the subject tokens a *prima facie* seeming that the capital is St. Louis and a *prima facie* seeming the capital is Jefferson City. Knowing that the former seeming is misleading, however, she judges that the capital is Jefferson City. Again, the *prima facie* seemings are FORs produced as the subject tries to remember the capital. However, on this picture, FORs are not the sole determiners of the doxastic attitudes that subjects form. It is an important component of the view that a weak FOR can spur the subject to engage in Type-2 reasoning that potentially reveals that the content of the FOR is false or misleading. Thus, it is consistent with the view that the FOR is defeasible, as in the Apparent Memory Case.

So feelings of rightness meet all of the necessary conditions on seemings, and we can understand the subjects of our central cases as tokening feelings of rightness. I

suggest, then, that all feelings of rightness are seemings, and moreover, that the seemings tokened in our central cases are all feelings of rightness. This view is a version of the Genus Account of seemings. Thus, it is compatible with mental states other than FORs counting as seemings, so long as these satisfy the necessary conditions delimiting the genus.<sup>39</sup> FORs, however, are special because they are not tied to any one source of belief. One can token an FOR while engaged in a variety Type-1 processes, so FORs are poised to shape our doxastic attitudes in a wide variety of circumstances.

## 6. Conclusions

I've defended the epistemic feeling account of seemings. Seemings are a genus of mental states that share certain important features: propositional content, a mind-to-world direction of fit, defeasibility, and presentational phenomenal character. Moreover, they satisfy the incompatible content norm. I've argued that we can better explain our central cases by supposing that subjects token *prima facie* seemings than by merely appealing to their dispositions to believe. Feelings of rightness—which help explain when subjects are more likely to simply endorse the output of some Type-1 process in judgement without engaging in Type-2 reasoning—satisfy all the necessary conditions on seemings and are a good fit for the mental states tokened by the subjects in our central cases. Thus, all token

---

<sup>39</sup> Perceptual experience, for example, can count as a seeming on those views according to which it is a propositional attitude with a mind-to-world direction of fit. Much of the debate over whether perceptual experiences should count as seemings concerns whether perceptual experiences have propositional content. For discussion, see Ghijzen (2015), Tucker (2010), Cullison (2013), and Bergmann (2013a, 2013b). I think, moreover, that cases of known illusion plausibly illustrate that perceptual experience satisfies the incompatible content condition. In the Müller-Lyer, my perceptual experience represents that the lines are different lengths, but I token an FOR that the lines are the same length when I remember how the illusion works.

feelings of rightness are token *prima facie* seemings, and the token *prima facie* seemings in our central cases are token feelings of rightness. On the epistemic feeling view, *prima facie* seemings play an important role in determining our doxastic attitudes.

I developed the epistemic feeling account of seemings without relying on linguistic intuitions about *seems* expressions. Moreover, the account did not depend on assuming that phenomenal conservatism is true, or that, more generally, the right view of justification appeals to seemings. That the epistemic feeling view is independent of these two issues is a significant advantage of the view.

### Chapter 3

#### Giving Liberals a Hand: A Defense of Liberalism Against Three Bayesian Arguments

Liberalism is the view that subjects can have immediate perceptual justification to believe at least some propositions about the external world, e.g., that *here is a hand*. Arguments in White (2006), Hawthorne (2004), and Cohen (2005) suggest that liberalism is incompatible or in tension with standard Bayesian confirmation theory.<sup>40</sup> Since Bayesian confirmation theory is such a fruitful paradigm, we have some reason to reject liberalism if the views are incompatible or in tension with one another.

A promising line of response to this argument, exemplified in Moretti (2015), is to argue that the views are not really in tension with one another, since the Bayesian argument against liberalism relies on several principles, which the liberal can reject, about the role that *beliefs about experiences* play in perceptual justification.<sup>41</sup> In this paper, I strengthen this line of response by reconstructing two ways in which the Bayesian argument could be revised so as not to depend on those controversial principles, and arguing that the liberal has a plausible response to each one. An important upshot, which hasn't been recognized in this literature so far, is that one of the central motivations for the Bayesian case against liberalism is a set of intuitions that has been interpreted as bearing on the rational force of perceptual experience, but which is more

---

<sup>40</sup> By standard Bayesianism, I'll mean the view that rational credence functions satisfy Kolmogorov's axioms (i.e., form a probability distribution) and are updated by strict conditionalization.

<sup>41</sup> Another line of response has been to revise Bayesian confirmation theory to better model a liberal view of perceptual justification. For this strategy, see Pryor (ms), Weatherson (2007), and Kung (2010).

plausibly interpreted as bearing on the rational force of *higher-order evidence*.

Recognizing the role that higher-order evidence plays in driving these intuitions will make clear that Bayesian confirmation theory gives us no reason to reject liberalism.

Here's the plan for the paper. In the first section, I'll explain in more detail what liberalism is and distinguish liberalism from several closely related views. In the second section, I'll present the first Bayesian argument against liberalism. In the third section, I'll explain why the argument depends on the assumption that perceptual justification is mediated by beliefs about experiences. In Sections 4 and 5, I reconstruct two ways of revising the Bayesian argument so that it does not depend on that assumption and argue that neither succeeds.

## 1. What is Liberalism?

Liberalism and conservatism are views about the structure of perceptual justification. Here's a rough gloss of the distinction.

**Liberalism:** Subjects can have immediate perceptual justification to believe certain propositions about the external world.<sup>42</sup>

**Conservatism:** Subjects can only have mediate perceptual justification to believe propositions about the external world. In particular, subjects have perceptual justification to believe propositions about the external world in virtue of having

---

<sup>42</sup> See Alston (1989), Audi (1993), Burge (2003), Huemer (2007), Peacocke (2004), and Pryor (2000).

independent justification (or default entitlement) to believe the denial of skeptical hypotheses.<sup>43</sup>

We'll also need a gloss of the distinction between mediate and immediate justification. It will be easiest to conduct the discussion here in terms of propositional rather than doxastic justification. Moreover, we'll be interested in what is required to have *some degree* of justification rather than knowledge-level justification. Here, then, is the gloss we'll use:<sup>44</sup>

**Mediate Justification:** S has mediate justification to believe p just when S has justification to believe p in virtue of having justification to believe q.

**Immediate Justification:** S has immediate justification to believe p just when it's not the case that S has justification to believe p in virtue of having justification to believe q.

The crucial term here is "in virtue of." The thought is that when a subject has mediate justification to believe p, the best philosophical explanation of why the subject has that justification must appeal to the fact that she has justification for some distinct belief q. (For an extended discussion of the distinction between mediate and immediate justification, see Chapter 4). I might have mediate justification to believe that orange

---

<sup>43</sup> See Wright (2002, 2004, 2007) and Davies (2000, 2003).

<sup>44</sup> This rendering of the distinction is loosely adapted from Alston (1983), who is concerned with immediate knowledge rather than immediate justification.

trees don't grow in Alaska because I have justification to believe that orange trees are tropical plants. Part of what explains why I have justification to believe the one proposition is that I have justification to believe the other.

A central motivation for adopting this explanationist account of the distinction is the inadequacy of accounts like the following.

**Simple Mediacy:** S has mediate justification to believe p just in case there is some proposition q such that S has justification to believe p only if S has justification to believe q.

Silins (2007) offers several plausible counterexamples to Simple Mediacy. For example, I might have justification to believe that the apple is red only if I have justification to believe that I exist, because if I have justification to believe *anything*, then I have justification to believe I exist. Still, it is implausible that my justification to believe the apple is red depends in the appropriate way on that other justification. The appeal to explanation avoids this difficulty. My justification to believe p is mediate just when the belief's status must be explained by (rather than merely enabled by or accompanied by) my having justification for some distinct belief.

So according to conservatism, subjects can only have mediate justification to believe propositions about the external world. In particular, subjects have perceptual justification to believe external world propositions partly in virtue of having independent justification (or entitlement) to deny skeptical hypotheses. Following Pryor (2004), in this context we can understand skeptical hypotheses as "non-perceiving hypotheses," or

hypotheses that are incompatible with the subject's perceiving her environment, e.g., that I'm BIV who merely hallucinates my hands or that I'm subject to an illusion. A subject's justification to deny a skeptical hypothesis is *independent* of her perceptual justification to believe p just when she could have that anti-skeptical justification even if she lacked perceptual justification to believe p (cf. Silins 2007, 110).

We should mention two points by way of explaining conservatism. First, some conservatives (but most prominently Wright (2004)) distinguish justification to believe a proposition from default entitlement to accept/reject propositions, and suggest that a subject's entitlement to reject skeptical hypotheses helps explain why she has perceptual justification to believe external world propositions. For example, on Wright's (2004) account, entitlement is distinct from justification in that:

[If] I am entitled to accept P, then my doing so is beyond rational reproach even though I can point to no cognitive accomplishment in my life, whether empirical or *a priori*, inferential or non-inferential, whose upshot could reasonably be contended to be that I...had succeeded in getting evidence justifying P. (174-5)

A version of conservatism that requires subjects to possess default entitlement to reject skeptical hypotheses is significantly less demanding than one that requires subjects to possess justification to disbelieve skeptical hypotheses. The Bayesian argument against liberalism, however, doesn't depend on the distinction, so for simplicity I will present the argument in terms of justification.

Second, as Pryor (2004) points out, we can define a broader version of conservatism that holds that a subject's perceptual justification to believe external world proposition p depends on her having independent justification to deny any undermining defeater for p. Undermining defeaters for p will include not only non-perceiving

hypotheses, but some additional kinds of propositions. For example, *that my experiences are generally unreliable* is not a non-perceiving hypothesis because it is compatible with my experience's constituting an instance of perception on this occasion. On this broader version of conservatism, a subject's perceptual justification for external world propositions depends on her having independent justification to deny these undermining defeater hypotheses as well. Whether we opt for a broader or narrower version of conservatism shouldn't matter for our purposes, so I will focus on the narrower version, according to which a subject's perceptual justification is taken to depend on her justification to deny just non-perceiving hypotheses.

Now that we have a better sense of what liberalism and conservatism hold, let's distinguish liberalism from some closely related views. Since liberalism holds only that subjects can be immediately perceptually justified in believing external world propositions, it doesn't give a complete theory of the factors in virtue of which experience has that justificatory power. Thus, liberalism is compatible with both access internalist and externalist views about perceptual justification, including process and virtue reliabilism, evidentialism, and phenomenal conservatism. It should be clear, then, that liberalism is distinct from phenomenal conservatism:

**Phenomenal Conservatism:** If it seems to S that p, then S has immediate *prima facie* justification to believe that p.

Phenomenal conservatism entails that undergoing a seeming as if p is sufficient for *prima facie* justification for belief that p. While phenomenal conservatism entails liberalism,

since it holds that certain beliefs about the external world are immediately justified by experience, liberalism doesn't entail phenomenal conservatism, since other views of perceptual justification can allow that experiences (under the right circumstances) immediately justify external world beliefs.

Liberalism is also compatible with both disjunctivist and common factor views about the structure of perception. On the latter, cases of perception and hallucination share some common factor, (e.g., perceptual experience), and cases of perception can be analyzed in terms of that common factor plus some appropriate causal connection between that factor and the external world. Analyzing perception and hallucination in terms of some common factor doesn't require that one take any particular stance on the possibility of immediate perceptual justification of external world beliefs. So liberalism and common factor views are compatible. On disjunctivism, cases of perception and hallucination share no common factor in terms of which they can both be explained; rather, perception and hallucination are distinct metaphysical kinds. Disjunctivists, too, can allow that external world beliefs can be immediately perceptually justified in both the good and bad cases, since allowing for the possibility of immediate perceptual justification does not require us to posit some common factor in terms of which both perception and hallucination are explained.

Finally, it is important to distinguish liberalism, a view about perceptual justification, from Mooreanism, a view about anti-skeptical justification. Consider the following Moorean inference:

P1. I have hands.

P2. If I have hands, then I'm not a BIV.

C. So I'm not a BIV.

Most liberals will accept that I can be immediately perceptually justified in believing P1. I needn't have independent justification to believe that my perceptual faculties are reliable, or that I'm not a BIV, in order for my experience as of my hands to justify me in believing that I have hands.<sup>45</sup> The Moorean goes a step further and holds that, not only can I be immediately perceptually justified in believing P1, but I can also become justified in believing the conclusion by competently performing the inference. While Mooreanism entails liberalism, liberalism doesn't entail Mooreanism. There are at least two ways to be a liberal while rejecting Mooreanism: One can deny that justification is closed under single-premise competent deduction, or one can accept justification closure but deny that the Moorean inference transmits warrant. On both options, a subject could be immediately perceptually justified in believing P1, competently perform the Moorean inference, yet fail to become justified in believing its conclusion.<sup>46</sup>

---

<sup>45</sup> Some liberals will deny that *I have hands* is the sort of proposition an experience could immediately justify one in believing. Some will hold, rather, that experiences only ever immediately justify subjects in believing propositions about low-level observable properties, such as shape, size, and color properties. In that case, we could imagine a Moorean inference whose first premise attributes one of these low-level observable properties.

<sup>46</sup> It's fairly unorthodox to accept liberalism but deny that the Moorean inference transmits warrant, but see Silins (2007) for a defense of this view. Most who have denied that the Moorean inference transmits warrant have appealed to conservatism to explain why transmission fails: Since a subject must have independent justification to believe P3 in order to be justified by experience in believing P1, she cannot become justified in believing P3 for the first time by competently performing the Moorean inference.

Mooreanism, liberalism, and phenomenal conservatism are each sometimes called *dogmatism*. Since these three views are distinct, I'll try to avoid confusion by avoiding the term *dogmatism* altogether.

## 2. The Bayesian Argument Against Liberalism

Several authors have argued that liberalism and standard Bayesianism are incompatible or at least in serious tension with one another.<sup>47</sup> Since Bayesianism is such a successful paradigm, this conflict seems to give us a good reason to reject liberalism. In this section, I'll explain the problem for liberalism before showing, in the next section, how the liberal can avoid it.

To see the problem, start by considering the following propositions.

E: I seem to have a hand.

H: I have a hand.

SK: I'm a BIV who merely seems to have a hand.

That I'm a BIV who merely seems to have a hand entails that I seem to have a hand: SK entails E. Since SK entails E, the following is true:

$$1. \quad \text{Cr}(E|SK) = 1$$

By Bayes' theorem:

---

<sup>47</sup> See, e.g., White (2006), Hawthorne (2004), and Cohen (2005). In presenting the argument, I'll follow White (2006) fairly closely. White explicitly targets a view that is closest to what I have called Mooreanism. However, the problem he identifies can be easily adapted into a problem against liberalism.

$$2. \text{Cr}(\text{SK}|\text{E}) = \frac{\text{Cr}(\text{E}|\text{SK})\text{Cr}(\text{SK})}{\text{Cr}(\text{E})}$$

From 1 and 2, it follows that:

$$3. \text{Cr}(\text{SK}|\text{E}) = \frac{\text{Cr}(\text{SK})}{\text{Cr}(\text{E})}$$

Since we shouldn't be antecedently certain either that SK is false or that E is true:

$$4. \text{Cr}(\text{SK}) > 0$$

$$5. \text{Cr}(\text{E}) < 1$$

It follows from 3, 4, and 5 that:

$$6. \text{Cr}(\text{SK}|\text{E}) > \text{Cr}(\text{SK})$$

$$7. \text{Cr}(\sim\text{SK}|\text{E}) < \text{Cr}(\sim\text{SK})$$

In other words, the rational response to learning E—that I seem to have hands—is to increase my rational credence in skeptical hypotheses that entail E—e.g., that I'm a BIV who merely seems to have hands. When I learn that E, my rational credence in the denial of skeptical hypotheses that entail E should decrease.

Finally, notice that H entails  $\sim\text{SK}$ : That I do have a hand entails that I'm not a BIV who merely seems to have a hand. Since H entails  $\sim\text{SK}$ , it follows that one's rational credence in  $\sim\text{SK}$  after learning E should be at least as great as one's rational credence in H after learning E. In other words:

$$8. \text{Cr}(\text{H}|\text{E}) \leq \text{Cr}(\sim\text{SK}|\text{E})$$

It's at this point that the problem for the liberal emerges. It follows from 7 and 8 that:

$$9. \text{Cr}(\text{H}|\text{E}) < \text{Cr}(\sim\text{SK})$$

One rationally must be *more* confident in the denial of the skeptical hypothesis than one is confident that one has hands after learning that it seems that one has hands. In other

words, if a subject is rationally confident that she has hands, given that it seems she has hands, then she is rationally confident to an even greater degree that she is not a BIV who merely seems to have hands.

There are at least two ways of understanding why 9 presents a problem for liberalism. On the first understanding, the problem is that liberalism struggles to explain why 9 is true, while conservatism offers a natural explanation. Consider a version of conservatism according to which subjects are perceptually justified in believing propositions about the external world in virtue of having independent justification (or default entitlement) to believe  $\sim$ SK. The thought is that this version of conservatism can explain 9: The reason why  $\text{Cr}(H|E) < \text{Cr}(\sim\text{SK})$  is that one must already have independent justification (or default entitlement) that renders one justifiably confident that one is not a BIV, in order for an experience to render one justifiably confident that one has hands. Liberalism has no comparable explanation for 9 available.

There is another way of understanding why 9 poses a problem for liberalism, which appeals more explicitly to bridge principles linking outright justification with rational credence. To see this alternative construction of the problem, note that it follows from 9 that:

10. For any threshold  $t$ , if  $\text{Cr}(H|E) > t$ , then  $\text{Cr}(\sim\text{SK}) > t$ .

Much more controversial is the following principle linking outright justification and rational credence.

11. There is some threshold  $t$  such that if  $\text{Cr}(p) > t$ , then  $S$  has justification to believe  $p$ .

It follows from 10 and 11 that:

12. If S has justification to believe H after learning E, then S has justification to believe  $\sim$ SK.

In other words, a subject who undergoes an experience as of her hands can be justified in believing that she has hands only if she has justification to believe that she is not a BIV. Liberalism actually is compatible with 12, since 12 holds only that it is a *necessary condition* for having justification to believe some external world proposition through experience that one have justification to believe the denial of the skeptical hypothesis. Liberalism denies only that one's anti-skeptical justification *helps explain* one's justification to believe external world propositions. Still, it might seem as though the conservative is much better positioned to explain 12 than the liberal. If liberalism is true, it is unclear why 12 should hold.

This reconstruction of the argument may be significantly less appealing to conservatives than the first, since it requires conservatives to defend the controversial 11. The principle is controversial because, among other reasons, it helps generate a version of the lottery paradox.<sup>48</sup> However, even if the liberal grants 11, she has available a strategy for resisting the argument that will apply equally well to either way of reconstructing the problem. I'll turn to this strategy in the next section.

Before doing that, we should briefly reply to one objection. One might worry that the Bayesian argument against liberalism is uninteresting because it depends on construing the skeptical hypothesis as *entailing* a subject's perceptual evidence. This might seem initially to be an artificial assumption designed to create a problem for the liberal.<sup>49</sup> However, we will still be able to derive the crucial inequality that purportedly

---

<sup>48</sup> See, e.g., Kyberg (1961).

<sup>49</sup> Thanks to Nic Koziol for raising this concern.

motivates conservatism (i.e., 9) even if SK does not entail E, so long as  $\text{Cr}(E|\text{SK})=1$ .

Moreover, even if the argument did depend on construing SK as entailing E, it would still represent a serious problem for the liberal, as we can easily generate skeptical hypotheses that do entail propositions about how things appear to the subject. That these hypotheses wouldn't exhaust the range of skeptical hypotheses does not lessen the seriousness of the problem.

### 3. Rejecting the Content and Belief Assumptions

The Bayesian argument against liberalism depends on certain assumptions about how to model rational belief change in response to experience. In particular, a subject who undergoes perceptual experience *e* is modeled, in the argument, as learning E, a proposition about how things appear to her. For example, a subject who has a visual experience as of her hands is modeled as updating by strict conditionalization on the proposition that *it seems that I have hands*. Thus, she assigns credence 1 to (and plausibly, forms an outright belief in) the proposition *it seems that I have hands*.

However, this picture makes two assumptions that the liberal can reject.<sup>50</sup> First, the liberal can reject:

---

<sup>50</sup> Here I'm closely following helpful discussion in Moretti (2015). While Moretti agrees that the liberal can reject the content and belief assumptions, he argues that she cannot do so without severely limiting the scope of beliefs that can count as immediately perceptually justified. I'll take up Moretti's view in the next section and argue that the liberal can reject the content and belief assumptions without this concession.

**Content Assumption:** When a subject undergoes a perceptual experience, she should be modeled as learning E, a proposition about her experience or about how things seem to her, rather than a proposition about the external world.

On liberalism, a subject can have perceptual justification to believe external world propositions even if she does not possess the introspective capacities required to form beliefs about her own experiences. Rather, on a liberal view, perceptual experiences are understood as informing subjects directly about the external world.

The liberal can also reject the following, closely related assumption:

**Belief Assumption:** Subjects acquire perceptual justification in virtue of updating by strict conditionalization on and assigning credence 1 to propositions.

The liberal can deny the belief assumption by holding that, so long as other conditions for justification are in place, a perceptual experience itself can justify beliefs about the external world. On a liberal view, a subject can have perceptual justification to believe propositions about the external world even before she assigns credence 1 to (and forms an outright belief in) any proposition in response to experience.<sup>51</sup>

The Bayesian argument depends crucially on the content and belief assumptions. By denying the content and belief assumptions, the liberal can deny that a subject who undergoes a visual experience as of her hands should be modeled as updating by strict conditionalization on E: *it seems that I have hands*. Rather, the liberal can hold that the

---

<sup>51</sup> This presupposes (safely, I think) that experiences are not a species of belief.

subject acquires, just in virtue of undergoing the experience, immediate perceptual justification to believe external world proposition H: *I have hands*. However, if we model the subject as learning H rather E, the Bayesian argument does not go through. Since H entails  $\sim$ SK, the following is true:

$$13. \text{Cr}(\sim\text{SK}|\text{H}) = 1$$

That is to say, on this way of modeling the case, the rational subject becomes certain that she is not a BIV upon having an experience as of her hands and learning H. It should be clear that 13 does not motivate conservatism about perceptual justification. We needn't suppose, to explain 13, that subjects must have independent justification (or default entitlement) to believe  $\sim$ SK before their perceptual experiences can justify them in believing external world propositions. Thus, the Bayesian argument against liberalism doesn't go through if we dispense with the content and belief assumptions.

#### **4. Revising the Bayesian Argument**

White (2006) recognizes that the liberal can respond to the Bayesian argument by rejecting the content and belief assumptions, but holds that a closely related version of the argument can go through without either assumption. In this section, I reconstruct this version of the argument and suggest that it depends on another controversial principle, which the liberal can reject, about how best to model the justificatory power of perceptual experience. I then identify an advantage of this line of response over a competing response in the literature.

According to the liberal, a subject can acquire perceptual justification to believe propositions about the external world even if she doesn't introspect on her experience and

form a belief about how things seem to her. However, as White points out, a suitably sophisticated subject *could* introspect on her experience. A reflective adult human who undergoes an experience *e* as of her hands could introspect on the experience and form the belief *E*, that it seems to her that she has hands. White suggests that “in this case the Bayesian argument goes through and seems unavoidable” (2006, 535). The thought here is that a sophisticated subject who both undergoes experience *e* and happens to form belief *E* should increase her credence in *SK*, since *SK* entails *E*, and that this will be enough to run the original Bayesian argument.

White continues by arguing that the justificatory power of a perceptual experience should not depend not whether a subject *happens* to form a belief *E* about her experience.

He writes:

If the rational response to its appearing that this is a hand, when I also believe that it appears that this is a hand, is to...[increase my confidence in *SK*], then surely this is the rational response to the same experience when I do not even consider how things appear to me. (2006, 535)

Here’s a reconstruction of the argument:

14. A sophisticated subject who undergoes an experience *e* as of her hands, and forms an introspective belief *E*—that it seems to her that she has hands—should increase her credence in *SK*—that she is a BIV who merely seems to have hands.

15. If 14 is true, then a subject who undergoes experience *e* without forming belief *E* should increase her credence in *SK*.

16. A subject who undergoes experience  $e$  without forming belief  $E$  should increase her credence in  $SK$ .

So a rational subject who merely has an experience as of her hands, without forming a belief about it, should respond to the experience by increasing her credence in the skeptical hypothesis.

If that's right, however, then the conservative can run a version of the Bayesian argument that does not depend on the content and belief assumptions. Here's how the revised Bayesian argument would go. Consider a subject who merely undergoes experience  $e$  and doesn't form introspective belief  $E$ . Suppose that  $t_1$  occurs before the subject undergoes an experience of her hands, and  $t_2$  afterward. If White is correct that such a subject should respond to the experience by increasing her credence in the skeptical hypothesis, then the following will be true:

$$17. Cr_{t_2}(\sim SK) < Cr_{t_1}(\sim SK)^{52}$$

Since  $H$  entails  $\sim SK$ , it will also be true that:

$$18. Cr_{t_2}(H) \leq Cr_{t_2}(\sim SK)$$

From 17 and 18, it follows that:

$$19. Cr_{t_2}(H) < Cr_{t_1}(\sim SK)$$

In other words, if an experience renders a subject rationally confident that she has hands, then she is rationally confident to an even greater degree that she is not a BIV. As before, the conservative seems to be in a better position to explain the inequality than the liberal,

---

<sup>52</sup> There is a quick way for the liberal to reject the argument here. Since we are assuming standard Bayesianism, the subject is updating by strict conditionalization on some proposition at this step. If the proposition is  $H$ , then 17 is false, and the argument doesn't go through. If the proposition is  $E$ , then the argument makes the Content Assumption, which the liberal rejects.

since on the conservative view, experiences only justify external world propositions for subjects who have independent justification (or default entitlement) to deny skeptical hypotheses. Notice, however, that this version of the Bayesian argument doesn't require subjects to form introspective beliefs about their experiences. Since the revised argument doesn't rely on the content and belief assumptions, the liberal needs some other response here.

Fortunately, the liberal has a plausible line of response here too. The liberal can reject 14: A rational subject who both undergoes experience  $e$  and forms belief  $E$  needn't increase her credence that she is a BIV. What motivates 14 is, in part, the thought that such a subject should be modeled as learning *just*  $E$ . After all, since  $SK$  entails  $E$ , it's true that  $Cr(SK|E) > Cr(SK)$ . However, this way of modeling the subject doesn't fully capture what the liberal takes to be the justificatory force of perceptual experience. On a liberal view, a subject who undergoes experience  $e$  and forms belief  $E$  already has propositional justification (so long as other conditions for justification are satisfied) to believe  $H$ . If we model the subject as learning just  $E$ , we don't take account of the justificatory force that, on a liberal view, accrues to the experience itself.

For that reason, it might seem that, on a liberal view, a subject who both undergoes experience  $e$  and forms belief  $E$  will best be modeled as learning both  $H$  and  $E$ . In that case, however, 14 is no longer plausible. (Recall that 14 holds that a sophisticated subject who undergoes an experience  $e$  as of her hands and forms an introspective belief  $E$  should increase her credence in  $SK$ .) If we model our subject as learning both  $H$  and  $E$ , then she should actually become less confident that  $SK$ .

$$20. Cr(SK|H \wedge E) < Cr(SK)$$

So 14 is no longer plausible if we model the subject in a way that more accurately reflects a liberal view of perceptual justification. Without 14, moreover, the conservative won't be able to show that the rational response to merely undergoing experience *e* (without forming the belief) is to become more confident in the skeptical hypothesis. In that case, the revised Bayesian argument won't go through.

Finally, it will be helpful to highlight one of the advantages of defending liberalism along these lines over an alternative proposal in the literature. Moretti (2015) responds to the revised Bayesian argument by accepting 14, but rejecting 15. So on this view 14 is true: The rational response to both undergoing *e* and forming belief *E* is to become more confident that *SK*. However, 15 is false: It doesn't follow from 14 that the rational response to simply undergoing *e* is to become more confident that *SK*. For Moretti, introspective beliefs about experiences have special justificatory force; they screen off the justificatory force of the experiences themselves. So when a subject both undergoes experience *e* and forms introspective belief *E*, that belief screens off the justificatory force of the experience: What one is justified in believing is then determined just by belief *E*.<sup>53</sup> For that reason, what a subject has justification to believe when she merely undergoes an experience can be different than what she has justification to believe

---

<sup>53</sup> The "screening off" view is motivated by a range of cases in which it seems intuitive that introspective beliefs screen off the justificatory force of the corresponding experiences. See, e.g., Moretti (2015, 274-5). At certain points, Moretti suggests that simply forming the relevant introspective belief is sufficient for it to screen off the justificatory force of the experience, writing, "Suppose...that *S* entertains the belief that *e*. As a consequence, the evidential force of *S*'s sole experience ... will be overruled by the evidential force of that introspective belief of *S*" (277). At other points, he suggests that screening off only happens under special circumstances, e.g, when *S* can determine the "evidential import" of her introspective belief (274). I'll consider the cases that motivate the *screening off* view in the next section, and argue that they are best interpreted as relying on intuitions about the rational force of *higher-order evidence*, rather than introspective belief per se.

when she both undergoes an experience and forms an introspective belief about it. On this view, a subject who both undergoes *e* and forms belief *E*—that it seems she has hands—should become more confident in SK, since her introspective belief alone has justificatory force. However, that doesn't mean that a rational subject who merely undergoes experience *e* must also become more confident that SK, since no screening off takes place in that case.

However, this response to the Bayesian argument has some unwelcome consequences for the liberal. On Moretti's view, when we form introspective beliefs about how things seem to us, what we are justified in believing is then determined by those beliefs, and not by the experiences themselves. So when a sophisticated subject introspects on her experiences, and forms a belief about how things seem to her, she is no longer immediately justified in believing anything on the basis of the experience. Rather, she is merely mediately justified by her belief about the experience. On this picture, immediate perceptual justification is fragile: It disappears when a subject introspects on her experience and forms a belief about how things seem to her. Moretti recognizes this consequence and argues that it shows that liberalism doesn't have much "anti-skeptical bite" (2015, 277).

However, as I've argued, the liberal isn't forced into this position: She has a motivated way of rejecting 14. On this way of defending of liberalism, immediate perceptual justification is robust: Subjects can be immediately perceptually justified in believing propositions even if they happen to introspect and form beliefs about how things seem to them perceptually. Since this response to the Bayesian argument preserves

the robustness of immediate perceptual justification, liberals should favor it over the alternative on offer.

## 5. Higher-Order Evidence and the Bayesian Argument from Cases

In this section, we'll reconstruct a second way of running the Bayesian argument against liberalism without the content and belief assumptions. On this way of running the argument, it is intuitive, in certain cases, that the rational response to perceptual experience is to increase one's credence in a skeptical hypothesis. But, so the argument goes, these cases are analogous to cases the liberal takes to be paradigmatic of immediate perceptual justification. The analogy is taken to support conservatism. I'll argue that the liberal has multiple lines of response available, depending on her stance on the rational force of higher-order evidence.

There seem to be many cases in which, intuitively, a subject rationally ought to increase her credence in a skeptical hypothesis in response to experience. Here is one such case, adapted from White (2006).<sup>54</sup>

**Mad Scientist:** There are ten slips of paper in a hat, numbered 1-10. A mad scientist will put Ann under anesthesia and draw a slip from the hat. If he draws 1-5, he'll cut off Ann's hands and replace them with fake hands, which are indistinguishable from her real hands. If he draws 6, he'll cut off her hands and leave the stumps. And if he draws 7-10, he'll leave her hands alone. Ann knows

---

<sup>54</sup> Wright (2007) and Vogel (2008) discuss similar cases.

all of this. Suppose that when Ann wakes up, she has a visual experience as of her hands. In fact, she has kept her real hands.

Consider the following propositions:

H: I have hands.

SK: I have fake hands.

At  $t_1$ , before Ann looks at the end of her arms, plausibly she should have  $Cr_{t_1}(H)=4/10$  and  $Cr_{t_1}(SK)=1/2$ . After undergoing a visual experience as of hands, Ann comes to know that the mad scientist didn't cut off her hands and leave the stumps. And so intuitively, she should become more confident both that she has hands *and that she has fake hands*. Plausibly, her credence at  $t_2$ , after undergoing the experience, should be  $Cr_{t_2}(H)=4/9$  and  $Cr_{t_2}(SK)=5/9$ . In other words, her experience intuitively should make her more confident that the skeptical hypothesis is true.

But, the conservative continues, the Mad Scientist case is similar in all relevant respects to cases that are, for the liberal, paradigmatic of immediate perceptual justification. Consider, for example, the following case:

**Position of Innocence:** Suppose that Ben has no reason (either empirical or *a priori*) to believe or disbelieve SK—that he has fake hands. He looks at the ends of his arms and has a visual experience as of hands.

The liberal holds that Ben can be immediately perceptually justified in believing that he has hands (so long as other conditions for justification are satisfied), even if he has no justification for the denial of the skeptical hypothesis. However, the conservative argues that Position of Innocence is similar in epistemically relevant respects to Mad Scientist. For example, in Mad Scientist, before undergoing the experience, Ann has  $Cr_{tl}(SK)=1/2$ , while Ben, before undergoing the experience, has no information concerning SK.<sup>55</sup> So if Ann rationally ought to become more confident in the skeptical hypothesis, after undergoing an experience as of her hands, it might seem that Ben rationally ought to respond in the same way to his experience in Position of Innocence.

Here's another way to put the point. In Mad Scientist, Ann's experience as of her hands intuitively should make her more confident that she has fake hands. Though her experience also makes her more confident that she has hands, it doesn't justify her outright in believing that. By contrast, the liberal holds that Ben, in Position of Innocence, needn't respond to his experience by increasing his credence in SK. Indeed, for the liberal, his experience can justify him outright in believing that he has hands. It's unclear what difference between the cases could explain why it is intuitive that Ann rationally ought to respond to her experience by becoming more confident that SK, while Ben needn't do the same.

Here's a recap of the argument so far:

---

<sup>55</sup> If one goes in for a version of the Principle of Indifference, then one might think that Ben should have  $Cr_{tl}(SK)=1/2$ , since he starts out with no information concerning SK. However, the conservative's argument doesn't depend on holding this.

21. In Mad Scientist, Ann rationally ought to respond to experience  $e$  by becoming more confident that SK.

22. There's no relevant epistemic difference between Mad Scientist and Position of Innocence.

23. So in Position of Innocence, Ben rationally ought to respond to experience  $e$  by becoming more confident that SK.

If the argument is sound, then the liberal is in trouble. For if Ben rationally ought to respond to his experience by becoming more confident that SK, the conservative can again show that:

24.  $Cr_{12}(H) < Cr_{11}(\sim SK)$

In other words, Ben must antecedently be more confident in the falsity of the skeptical hypothesis than he is confident that he has hands, after undergoing the experience. But, just as before, the inequality motivates conservatism: the conservative can argue that the reason why 24 is true is that Ben must have independent justification to believe that SK is false, if his experience is to justify him in believing some proposition about the external world.

Notice that this Bayesian argument doesn't depend on either the content or belief assumptions. Nowhere is it assumed that Ann or Ben introspect on their experiences and form beliefs about how things seem to them. So the liberal cannot respond to this version of the Bayesian argument in the same way as she does to the original version.

The liberal might be tempted to respond by denying 22 and identifying a difference between the two cases that can explain why Ann and Ben should respond differently to their experiences. One clear difference between the cases is that Ann has an undermining defeater that Ben lacks. Since Ann knows the mad scientist's plan, she knows that taking her experience at face value isn't a reliable way of forming true beliefs. However, notice that we can still run the argument even if we amend the Mad Scientist case so that Ann no longer possesses an undermining defeater. For example, we can amend the case to make it very unlikely that the mad scientist gives Ann fake hands, and very likely that he leaves Ann's hands alone.

**Mad Scientist 2:** There are a thousand slips of paper in a hat, numbered 1-1000.

A mad scientist will put Ann under anesthesia and draw a slip from the hat. If he draws 1, he'll cut off Ann's hands and replace them with fake hands, which are indistinguishable from her real hands. If he draws 2, he'll cut off her hands and leave the stumps. And if he draws 3-1000, he'll leave her hands alone. Suppose Ann, knowing all this, wakes up and takes a look at the end of her arms. She has a visual experience as of hands. Suppose that, in fact, she has kept her real hands.

In this case, it's overwhelming likely that Ann will wake up with her hands untouched. And so plausibly she does not possess a traditional undermining defeater. Nevertheless, as before, it is intuitive that Ann rationally ought to respond to her experience by increasing her credence that she has fake hands. Before undergoing the experience, she should have  $Cr_{t1}(SK)=1/1000$ , but after the experience, she should have  $Cr_{t2}(SK)=1/999$ .

And this is enough to generate the problem for liberalism. So the liberal can't respond to the argument by holding that the difference between Mad Scientist and Position of Innocence, which explains why Ann and Ben should respond differently to their experiences, is that Ann has an undermining defeater that Ben lacks.

However, the liberal has another response available to her. There is an important difference between even the revised Mad Scientist case and Position of Innocence. In Mad Scientist 2, Ann possesses higher-order evidence that Ben lacks. In particular, Ann knows how reliable she'd be if she believed that she has hands on the basis of her experience: She'd believe H truly in that way with 998/999 reliability. In Position of Innocence, by contrast, Ben has no information about how reliable he'd be if he believed he has hands on the basis of his experience.

There's a great deal of controversy about whether (and how) higher-order evidence determines which first-order doxastic state is rational for a subject. It would be less than ideal if the liberal had to commit to a particular view on the force of higher-order evidence in order to respond to the conservative here. However, I want to suggest that the liberal has multiple ways of responding to the argument available to her, depending on the view she adopts about higher-order evidence. There is a way for the liberal to respond to the argument if she accepts that higher-order evidence helps determine which first-order doxastic state is rational for a subject, and there is a way for the liberal to respond if she denies this. So the liberal can respond to the argument while effectively remaining neutral on the higher-order evidence debate.

Before explaining these options, it will be helpful to highlight more explicitly the parallel between Mad Scientist 2 and paradigmatic cases in the higher-order evidence literature. Consider, for example, the following widely-discussed case:

**Hypoxia:** Cate is a pilot, trying to determine whether she has enough fuel to reach her destination. She performs some calculations and concludes that she does have enough fuel. Ground control then informs her that she's flying at an altitude that puts her at risk for hypoxia, which severely limits reasoning abilities. They let her know that pilots flying at her altitude who try to determine whether they have enough fuel using the calculations she performed arrive at the wrong answer half of the time.

Many have found it intuitive that Cate should adopt credence .5 that she has enough fuel to reach her destination, since she knows that subjects who perform calculations  $C$  at her altitude only form true beliefs with 50% reliability. Regardless of whether one thinks that such intuitions should be accepted or explained away, Ann in Mad Scientist 2 is relevantly like Cate. After waking up and having a visual experience as of her hands, Ann knows that were she to believe that she has fake hands, she would believe truly with  $1/999$  reliably. Moreover, just as many find it intuitive that Cate should adopt credence .5 that she has enough fuel to reach her destination, we find it intuitive that Ann should have  $C_{r2}(SK)=1/999$ .

The liberal can exploit the parallel between Hypoxia and Mad Scientist 2, regardless of the stance she takes in the higher-order evidence debate in general. Recall the conservative's argument from analogy:

21. In Mad Scientist 2, Ann rationally ought to respond to experience *e* by becoming more confident that SK.

22. There's no relevant epistemic difference between Mad Scientist 2 and Position of Innocence.

23. So in Position of Innocence, Ben rationally ought to respond to experience *e* by becoming more confident that SK.

Consider, first, the liberal who holds that a subject's higher-order evidence does help determine what first-order doxastic state is rational. Such a liberal is well-positioned to deny 22. The relevant epistemic difference between the cases is that Ann has higher-order evidence about her own reliability that Ben lacks, and this can explain why they rationally ought to respond differently to experiences as of their hands.

One way (but certainly not the only way) for the liberal to cash out this idea would be to embrace a calibrationist view about the significance of higher-order evidence about one's own reliability.

**Calibrationism:** If S knows that her degree of reliability with respect to p is r, independently of her reasoning on the first order, then S rationally ought to adopt credence r toward p.<sup>56</sup>

On this line, Ann knows that her degree of reliability, before undergoing the experience as of her hands is 998/1000 with respect to H, and 1/1000 with respect to SK. She also knows that, after having an experience as of her hands, her degree of reliability with respect to H increases to 998/999, and her degree of reliability with respect to SK increases to 1/999. By calibrationism, she rationally ought to assign  $Cr_{t1}(H)=998/1000$  and  $Cr_{t1}(SK)=1/1000$  before her experience, and to assign  $Cr_{t2}(H)=998/999$  and  $Cr_{t2}(SK)=1/999$  after her experience. That's why she rationally ought to respond to her experience as of a hand by, not only becoming more confident that she has hands, but also becoming more confident in the skeptical hypothesis.

In Position of Innocence, however, Ben lacks knowledge of his degree of reliability, and so does not satisfy calibrationism's antecedent. Thus, calibrationism doesn't constrain Ben's credences. After undergoing the experience as of his hands, Ben is free to raise his credence in H and lower his credence in SK.<sup>57</sup> It's this asymmetry between higher-order evidence between the two cases that explains why, while it's

---

<sup>56</sup> More sophisticated versions of this principle have been defended by Christensen (2010), Horowitz & Sliwa (2015), and White (2009). For a critique of calibrationism, see Schoenfield (2015).

<sup>57</sup> One might object in the following way. Ben's expected degree of reliability with respect to SK is .5; after all, Ben has no information about how reliable he would be if he would be if he believed SK. So when Ben has an experience as of his hands, he rationally ought to assign credence .5 to SK. Thus, both Ann and Ben must calibrate their credences, so there is no relevant epistemic difference between them.

However, Ben still fails to satisfy the antecedent of calibrationism. Even if his expected degree of reliability with respect to SK is .5, he plausibly still does not count as knowing this.

rational for Ann to respond to her experience by becoming more confident that she has fake hands, Ben is free to respond to his experience by becoming more confident that he has hands and less confident that he has fake hands.

Of course, there are many reasons why the liberal might be wary of calibrationism. It's important to emphasize that the liberal needn't embrace calibrationism as the best account of the rational force of higher-order evidence in order to respond to the conservative. It would be enough to argue that Ann has higher-order evidence that Ben lacks, and that this makes a relevant difference as to how it is rational for each of them to respond to their experiences. I mentioned calibrationism as one way of cashing out this strategy, for the sake of concreteness, but it is strictly more than the liberal needs.

Let's consider now a liberal who denies that a subject's higher-order evidence helps determine the first-order doxastic state it is rational for her to adopt. The most familiar view of this sort is a Right Reasons view, according to which the doxastic state it is rational for a subject to adopt is fully determined by her first-order evidence (and her evidence-evaluating capacities or competences).<sup>58</sup> On such a view, Cate's credence that she can safely make the trip should be determined just by the first-order calculations, while Ann's credence that she has hands should be determined just by her first-order perceptual evidence. This account of the rational force of higher-order evidence will open the door for the liberal to deny 21. On at least certain liberal views of perceptual evidence—according to which Ann's first-order perceptual evidence in Mad Scientist 2 is simply *that she has hands* and so is incompatible with the skeptical hypothesis—21 will

---

<sup>58</sup> See, e.g., Kelly's (2005) account of peer disagreement. He has since given up the Right Reasons account in favor of the Total Evidence account in Kelly (2010).

be false: It is not the case that Ann rationally ought to respond to her experience by becoming more confident in the skeptical hypothesis.

Level-splitting views of this sort will need an error theory for the intuition prompted by cases like Hypoxia and Mad Scientist 2. The thought would be that, while it might be intuitive that Ann in Mad Scientist 2 and Cate in Hypoxia rationally ought to adjust their credence in light of their knowledge of their own reliability, this intuition is misleading. For example, Schoenfield (2015) has argued that, though calibrationism is not a norm of rationality, we have calibrationist intuitions because calibrationism is a *principle of reasoning*.<sup>59</sup> While principles of rationality “take as input total bodies of evidence and output a belief state,” principles of reasoning are “principles about which transitions of thought, or reasoning processes, should occur in the process of deliberation” and “take as inputs beliefs, seemings, judgments, hunches, credences, and so forth and can output similar sorts of things” (435). Thus, our calibrationist intuitions in response to Hypoxia and Mad Scientist 2 are explicable in light of the fact that calibrationism is a principle of reasoning, even if not a principle of rationality. A full evaluation of this suggestion is beyond the scope of this paper. The important point is that, even if the liberal denies that higher-order evidence has any role to play in determining what doxastic state it is rational for a subject to adopt, she still has a strategy for responding to the Bayesian argument from cases: Namely, she can appeal to an error theory from the broader higher-order evidence literature for the calibrationist intuitions that motivate 21. Thus, the liberal can accept that the Bayesian argument from cases is

---

<sup>59</sup> Schoenfield (2015) does not advocate the Right Reasons view, but is concerned to explain the source of our calibrationist intuitions in light of several apparently serious problems facing calibrationist norms of rationality.

motivated by intuitions about higher-order evidence even if she does not give much weight to such intuitions.

Let's take stock. We first saw a version of the Bayesian argument against liberalism that doesn't depend on the belief and content assumptions, and so can't be dealt with by denying those assumptions. The worry was that there are certain cases in which it is intuitive that the rational response to perceptual experience is to increase one's credence in a skeptical hypothesis, but that these cases are similar in relevant respects to cases that the liberal takes to be paradigmatic of immediate perceptual justification. I've argued, in response, that the best explanation of why, in certain cases, the rational response to experience is intuitively to increase one's credence in the skeptical hypothesis is that subjects in these cases possess higher-order evidence about how reliable they'd be if they believed the skeptical hypothesis on the basis of their experience. I've suggested, moreover, that subjects lack this higher-order evidence in those cases that are paradigmatic of immediate perceptual justification for the liberal. Finally, I've argued that the liberal can pursue this response to the Bayesian argument no matter what stance she takes in the wider debate about the rational force of higher-order evidence.

## **6. Conclusions**

I've argued, in this paper, that the Bayesian argument against liberalism does not succeed and cannot be successfully revised. We saw, first, that the original version of the argument depends on the content and belief assumptions. The liberal can respond by denying that perceptual justification requires subjects to form beliefs about the character

of their experiences. Next, we considered one way to revise the Bayesian argument so that it no longer depends on those assumptions: The thought was that if the rational response to both undergoing experience *e* and forming belief *E* is to increase one's credence in the skeptical hypothesis, then this must be the rational response to simply undergoing experience *e*. However, the liberal has a plausible response here too: She can deny that the rational response to undergoing experience *e* and forming the belief *E* is to become more confident in the skeptical hypothesis. Finally, we considered a second way to revise the Bayesian argument so that it does not depend on the content and belief assumptions. The thought here is that it is intuitive, in certain cases, that the rational response to undergoing a perceptual experience is to increase one's confidence in the skeptical hypothesis, but that these cases are similar in relevant respects to cases that are supposed to be paradigmatic instances of immediate perceptual justification. However, the liberal can deny that the cases are relevantly similar by pointing out an asymmetry in the subjects' higher-order evidence, and she can pursue this line of response regardless of the stance she takes on the rational force of higher-order evidence in general. Thus, a commitment to standard Bayesianism does not force us to give up liberalism about perceptual justification.

## Chapter 4

### The Scope of Immediate Perceptual Justification

According to liberalism, at least some beliefs about the external world can be immediately perceptually justified. Roughly, a belief is immediately justified when it is justified in a way that does not require that it be based on other justified beliefs. Though liberalism is an attractive view, it faces a difficult challenge in delimiting the scope of immediate perceptual justification. While it seems plausible that I could be immediately perceptually justified in believing that *that's red*, for example, it seems far less plausible that I could be immediately perceptually justified in believing something like *this house was built in 1958*. Other cases are more difficult: It is unclear whether I could be immediately perceptually justified in believing a proposition like *that's Janet*, or *that's a dendrobium orchid*.

It is tempting to think that the contents of perceptual experience determine what we can be immediately perceptually justified in believing. I can be immediately perceptually justified in believing that *that's red*, but not that *the house was built in 1958*, because I could token a perceptual experience with the former but not the latter content. Moreover, it is controversial whether I could be immediately perceptually justified in believing that *that's Janet*, or that *that's a dendrobium orchid*, because it is controversial whether we could have perceptual experiences with contents like these.

In this paper, I'll argue that this natural thought is incorrect: The scope of immediate perceptual justification is not fully determined by the contents of perceptual

experience. It is possible for a subject, who tokens experience *e*, to be immediately perceptually justified in believing *p*, even though *p* is not among the contents of *e*. This gap in content, moreover, does not owe to any general difference in the types of content that experiences and beliefs can have, e.g., the difference between non-conceptual and conceptual content. In Part 1 of the paper, I'll discuss several cases from McGrath (2016) and Silins (2011) with this structure, which I'll call *mismatch cases*. I'll argue that they are in fact cases of immediate perceptual justification in which the content of the justified belief outruns, in a sense that I'll specify, the content of perceptual experience.

In Parts 2 and 3, I'll consider the upshot of the argument in Part 1 for liberalism. The possibility of mismatch cases tells against versions of liberalism that require a tight semantic connection between the contents of experience and the contents of the beliefs that they immediately justify. I'll consider three versions of phenomenal conservatism—which is sometimes expressed as the principle that *S* has *prima facie* immediate justification to believe *p* iff it seems to *S* that *p*—and argue that because they require this tight semantic connection, they cannot explain immediate justification in mismatch cases.

Finally, I'll sketch two general strategies for accounting for mismatch cases within the liberal camp. The *permissive evidentialist strategy* accounts for mismatch cases by adopting a more permissive notion of the evidential support relation that must obtain between the content of an experience and the belief that it immediately justifies. The *basic method strategy* handles mismatch cases by accounting for immediate justification in terms of non-inferential competences, methods, or processes, and denying that any particular relation must hold between the content of an experience and immediately

justified perceptual belief. At the end of the paper, I briefly argue that the basic method strategy is the more promising solution.

The argument in this paper owes a great deal to those in McGrath (2016) and Silins (2011). Nevertheless, I want to highlight the ways in which I depart from each of them. Silins would likely accept that the mismatch cases are instances of immediate perceptual justification. He uses mismatch cases to argue that we can be immediately justified in believing “indirect contents of experience,” i.e., contents that an experience has “at least in part in virtue of having some other content” (2011, 354). By contrast, I use the mismatch cases to argue for *basic method* views of justification, which allow for an even more radical divergence between the content of experience and the contents of immediately perceptually justified beliefs than Silins recognizes.

Unlike Silins and I, McGrath (2016) ultimately denies that the mismatch cases are instances of immediate justification. He uses the mismatch cases to argue that the most defensible liberal position is one that allows a wide range of beliefs with “high-level” content to count as immediately perceptually justified. I agree with this, but I argue that even versions of phenomenal conservatism that allow “high-level” beliefs to be immediately perceptually justified struggle to handle mismatch cases. Only certain liberal views, I suggest, can adequately capture them.

## **1. Mismatch Cases**

Consider the following cases, adapted from McGrath (2016) and Silins (2011).

**Case A—Negation:** Anita sees the sky, which is a striking shade of blue. In fact, the shade is indigo. Suppose that Anita is not quite sure what indigo is; she is about as confident that indigo is a shade of blue as a shade of yellow. She doesn't form the belief that the sky is indigo. However, Anita does know what turquoise is. When she forms beliefs about whether some perceived object is or isn't turquoise, her beliefs are reliably true. Suppose that Anita, looking at the sky, believes truly that the sky isn't turquoise (McGrath 2016, 116).

**Case B—Disjunction:** Ben doesn't have perfect pitch, but he can reliably recognize and name certain notes. When an A, C, or D is played on an instrument, he can reliably tell that it's one of those notes, but he cannot reliably tell which note it is. (That is to say, were he to believe that *the note is either an A, C, or D*, his belief would be reliably formed, but were he to believe that *the note is an A*, say, his belief would not be reliably formed.) Suppose that someone strikes a C on the piano, and Ben comes to believe truly that the note is either A, C, or D (McGrath 2016, 115).

**Case C—Comparative Similarity:** Clifton sees a tangerine paint swatch. Suppose Clifton is reliable at determining when a perceived color is more similar to yellow than purple. (Though *being more similar to yellow than purple* is a vague predicate, I'm assuming that there are some cases where it is clearly true or false that a color is more similar to yellow than purple. See Byrne (2005) for discussion.) Clifton forms the true belief that the color of the swatch is more similar to yellow than purple (Silins 2011, 350).

Part I of the paper is devoted to arguing for a particular interpretation of these cases. First, I will argue that each subject's belief is immediately justified. Second, I'll argue that the content of each subject's belief outruns the content of her experience, in a sense that I'll specify.<sup>60</sup>

### *1.1 Immediate vs. Mediate Justification*

It's crucial that we clarify the distinction between mediate and immediate justification. Our focus here will be on the distinction between immediate and mediate *doxastic* justification, rather than propositional justification. That is, we'll be concerned with whether a subject's belief that *p* is immediately or mediately justified, rather than whether a subject *has* immediate or mediate justification to believe that *p*.

The rough idea is that mediately justified beliefs depend for their justification on other beliefs of the subject, while immediately justified beliefs do not. The challenge is to specify what type of dependence on other beliefs is relevant. It will be helpful to start with a rendering of the distinction adapted from Alston (1983):

**Mediate Justification:** *S* is *prima facie* mediately justified in believing *p* iff *S*'s belief that *p* is *prima facie* justified "by virtue of some relation this belief has to some other justified belief(s) of *S*" (74).

---

<sup>60</sup> It may seem inapt to say that the content of the subjects' beliefs "outruns" the content of their experiences, because plausibly the content of the experiences entails the content of the respective beliefs. See Section 3 for discussion of the significance of these entailments. I mean "outruns" in the sense of "is true in a wider range of circumstances." So, e.g., Anita's belief that the sky isn't turquoise outruns the content of her visual experience as of the indigo sky, because it can be true that the sky isn't turquoise when the sky is a variety of colors besides indigo.

**Immediate Justification:** S is *prima facie* immediately justified in believing p iff S's belief that p is *prima facie* justified "by virtue of something other than some relation this belief has to some other justified belief(s) of S" (74).

To understand the distinction, we'll first want to know what "by virtue of" means here. The thought is that, to determine whether S's belief is mediately or immediately justified, we must turn to the best philosophical explanation of why S's belief is justified. If according to the best philosophical explanation, S's belief that p is justified because of some relation it bears to other of S's justified beliefs—e.g., being based on or inferred from other justified beliefs—then S is mediately justified in believing that p. By contrast, immediate justification is, as Alston puts it, a "wastebasket category," because a belief will count as immediately justified on a variety of explanations of its justification (1983, 75). For example, a belief will count as immediately justified if, according to the best philosophical explanation of its justification, the belief is justified because of the fact that makes it true, simply because it is believed or understood, because it is based on a reliably formed perceptual experience or intuition, because it is the output of a non-inferential competence or reliable process, etc. (1983, 75).<sup>61</sup>

Crucially, S's belief that p can be immediately justified compatibly with the belief's depending in certain ways on other justified beliefs of S. Suppose we assume (just

---

<sup>61</sup> Sometimes "inferential" is used in a very broad sense according to which any transition among mental states, even subpersonal ones, is an inference. Here, we're stipulatively using "inferential" in a narrower way, according to which the premises of the inference are supplied by beliefs. This sense of "inference" is also distinct from Siegel's (2017) view of inference as a "distinctive kind of response to an informational state ... that produces a conclusion," which does not require "reckoning" or "taking some information to support a conclusion" (77-78).

for illustration) a view of concept possession according to which S possesses the concepts that figure in her belief that p only if she has certain other justified beliefs.<sup>62</sup> For example, suppose that S counts as possessing the concept CAT only if she justifiably believes that cats are animals. In that case, S is justified in holding any belief in which the concept CAT figures only if she is justified in believing that cats are animals. For all that, S might still be immediately justified in believing that *that's a cat*. S's belief that *that's a cat* need not be inferred from or based on her justified belief that cats are animals. The best philosophical explanation of why S can token the belief that *that's a cat* might appeal to her justifiably believing that cats are animals. However, the best philosophical explanation of why S's belief that *that's a cat* is justified needn't appeal to that other belief. S's believing that cats are animals is an enabling condition on her justifiably believing that *that's a cat*, rather than the reason for which she believes that.<sup>63</sup>

It will be helpful to clarify briefly how the distinction above is related to the varieties of foundationalism. Foundationalist views are standardly thought to have two components. Here, e.g., is Lyons (2009): "Foundationalism is the view that (a) there is a privileged class of basic beliefs, that is, beliefs whose justification does not depend on inferential or evidential connections to other beliefs, and (b) all non-basic beliefs, if justified, ultimately derive their justification from evidential relations to these basic beliefs" (3). While (a) and (b) are a common package, (a) does not entail (b). Foundationalist views are committed to our having some immediately justified beliefs (a). However, one could accept that we have some immediately justified beliefs (a), while

---

<sup>62</sup> I'm not endorsing this view of concept possession, but rather using it to illustrate what immediate justification requires. We here go beyond formulations of theory-theory accounts of concepts on which a mental theory is a set of beliefs (whether justified or not).

<sup>63</sup> For discussion, see Alston (1983, 78-79).

denying that all mediately justified beliefs ultimately derive their justification from evidential relations to immediately justified beliefs (b).<sup>64</sup>

It should be clear that nothing in our definition of immediate justification requires that immediately justified beliefs be infallible, indubitable, or incorrigible. Thus, one can be committed to our having immediate justification in this sense without being committed to strong Cartesian foundationalism. However, the definition leaves open whether we are to understand “justification” as knowledge-level justification or some weaker epistemic status falling below knowledge-level justification. Thus, it will be helpful to distinguish between moderate and weak immediate justification.<sup>65</sup> S’s belief that *p* enjoys *moderate* immediate justification when it enjoys knowledge-level immediate justification. S’s belief enjoys *weak* immediate justification when it enjoys immediate justification falling short of that required for knowledge. The distinction between moderate and weak justification is orthogonal to the distinction between *prima facie* and *ultima facie* justification. Even knowledge-level immediate justification is defeasible (at least for the fallibilist views we’ll focus on here). For example, S’s moderate immediate justification to believe the wall is red might be defeated when she learns that a tricky red light is shining on the wall.

I’ll be arguing that the mismatch subjects enjoy moderate immediate justification. Perhaps surprisingly, the possibility of moderate immediate justification is compatible with certain views of justification that incorporate coherentist elements. Consider, for example, Goldberg’s (2012) “reliabilist foundationalist coherentism.” The view is

---

<sup>64</sup> For further discussion, see Pryor (2000, 535).

<sup>65</sup> This follows Bonjour’s (1985) classic distinction between modest and weak/minimal foundationalism.

committed to moderate immediate justification: Beliefs that are the outputs of reliable belief-independent processes, (i.e., any process that does not take beliefs as inputs), possess knowledge-level immediate justification. However, the reliability of some such processes owes to “coherence-monitoring filters” that “interrupt belief formation whenever an incoherence with background information is detected” (2012, 188). Similarly, Fleisher’s (2019) method coherence reliabilism holds that some beliefs possess knowledge-level immediate justification, but only if they are the outputs of a reliable method that has been “coherence-tested” against other methods.<sup>66</sup> I mention these views not to endorse them, but to illustrate the thought that a commitment to moderate immediate justification is compatible with views of doxastic justification that incorporate coherentist elements.

Finally, we should note that, at least for some propositions *p*, it is possible to be both mediately and immediately justified in believing *p*. For example, Ann might be immediately perceptually justified in believing the apple is red when she sees it. At the same time, Bob might be mediately justified in believing the apple is red by competently inferring that proposition from some other pieces of propositional knowledge, e.g., that Cate bought the apple and only buys red apples. Indeed, a single subject can have both mediate and immediate *propositional* justification to believe *p* at the same time, as when she both sees the apple and knows the premises of the inference just outlined. That a subject has immediate justification to believe *p* is compatible with her also having mediate justification to believe *p*.

---

<sup>66</sup> Roughly, two methods are coherence tested against one another when they output two “sufficiently similar” beliefs at the same time, where “sufficient similarity” is understood in terms of mutual probabilistic support.

## *1.2 The Content Constraint*

I'll be arguing that the mismatch subjects' beliefs are moderately immediately justified, i.e., they enjoy knowledge-level justification that does not depend (in the sense laid out above) on other justified beliefs of the subject. Next, let's clarify the sense in which the contents of the subjects' beliefs outrun the contents of their perceptual experiences. As a first pass, we might think of the cases as (putative) counterexamples to what Silins (2011) has called the content constraint:

**Silins's Content Constraint:** "If S's [perceptual] experience *e* gives S immediate justification to believe some external world proposition that *p*, then it is a content of *e* that *p*" (349).

On this picture, the mismatch cases are putative counterexamples to the content constraint because the contents of the subjects' immediately justified perceptual beliefs are not contents of their experiences. However, we might worry that, on many views of the content of perceptual experience, the content constraint will be trivially false. Consider, for example, views according to which experiences and beliefs have different kinds of content, e.g., that experiences have non-conceptual content while beliefs have conceptual content, or that experiences have accuracy conditions while beliefs have truth conditions. On these sorts of views, the content constraint will be trivially false because experiences and beliefs will never share content. However, I want to argue that, even on

views where beliefs and experiences have the same general type of content, there can be a gap between immediately justified perceptual belief and the content of experience. This gap is not merely a gap between conceptual and nonconceptual content, or between truth conditions and accuracy conditions, etc.

For these reasons, I'll focus instead on arguing against the following principle.

**Content Constraint:** If S's perceptual experience *e* gives S immediate justification to believe some external world proposition *p*, then *p* is true iff *e* is accurate.

First, notice what the Content Constraint is *not* saying: It does not hold that all immediately justified beliefs are true, or that only subjects in the good case enjoy immediate perceptual justification. Rather, Content Constraint captures the thought that there needs to be a tight semantic connection between the contents of immediately justified perceptual beliefs and the experiences that cause them, or on which they are based. However, Content Constraint captures this semantic connection without assuming that beliefs and experiences have the same general type of content. Moreover, it helps us avoid dealing with the tricky issue of how to think about entailment among contents that are different general types or formats.<sup>67</sup> Finally, the Content Constraint does presuppose that immediate perceptual justification requires that experiences have content. This seems to be a weakness of this formulation of the principle, since naive realist and adverbialist

---

<sup>67</sup> For Content Constraint to be plausible, we need to think of *e* narrowly as that aspect of S's full perceptual experience that immediately justifies her in believing *p*. If we take *e* to encompass S's full perceptual experience, then Content Constraint will be trivially false.

views of perception plausibly are compatible with the possibility of immediate perceptual justification. Since I'll be arguing against the Content Constraint, I'll set aside this complication here.

## 2. The Mismatch Argument

Let's turn now to arguing for the proposed interpretation of the mismatch cases. I'll argue for the following two theses:

1. In each case, the subject's belief that *p* is immediately justified.
2. For each subject's perceptual experience *e* and immediately justified belief that *p*, it's not the case that *p* is true iff *e* is accurate.

It will be more difficult to argue for the first thesis than the second, so let's start with the second.

### 2.1 *Content Mismatch*

It's plausible that the content of the subject's belief in each case outruns the content of the relevant experience. Anita's visual experience, though it represents the color of the sky, doesn't represent that the sky is not turquoise. Ben's auditory experience, though it represents the pitch of the note, doesn't represent that the note is

either A, C, or D. Finally, Clifton's visual experience represents the swatch's color, but it doesn't represent that the color is more similar to yellow than it is to purple.<sup>68</sup>

More carefully, for each subject's immediately justified belief that *p*, it's not the case that *p* is true iff *e* is accurate. Consider, for example, Case A. In this case, *p* is the proposition that the sky isn't turquoise, and *e* is Anita's visual experience as of the indigo sky. It is possible for *p* to be true while *e* isn't accurate: The sky can fail to be turquoise (*p*) compatibly with it's being—not indigo—but rather gray, yellow, or purple. (We're just focusing on the content of Anita's actual experience as of the indigo sky and asking whether that content must be accurate iff it is true that the sky isn't turquoise.) So it's not the case that, if the content of Anita's belief is true, then her experience is accurate.

The same is true of Case B. Here, *p* is the proposition that the note is either A, C, or D, and *e* is Ben's auditory experience as of the note, which is in fact C. Again, it's possible for *p* to be true while *e* isn't accurate: The note could be A, C, or D compatibly with it's being an A. However we understand the content of Ben's auditory experience, it should be clear that it is possible for that content to be inaccurate while it is true that the note is either A, C, or D.

Finally, Case C follows the same pattern. In case C, *p* is the proposition that the color is more similar to yellow than to purple, and *e* is Clifton's visual experience as of the tangerine swatch. It is possible for *p* to be true while *e* is inaccurate: There are colors besides tangerine that are more similar to yellow than purple, e.g., yellowish green. So here again, it is possible for the content of the subject's belief to be true, while his

---

<sup>68</sup> See for discussion McGrath (2016, 115-16) and Silins (2011, 350).

experience is inaccurate. Hence the “mismatch” between the content of the subject’s belief and the relevant aspect of experience.

To make this point, we don't need to rely on any particularly contentious thesis about the contents of perceptual experience. For example, we do not need to come down on whether liberalism or conservatism about experiential content is true. Here’s a rough gloss of that distinction.

**Liberalism:** Perceptual experiences can represent “high-level” properties such as natural and social kinds, emotions, causation, affordances, etc.<sup>69</sup>

**Conservatism:** Perceptual experiences represent only “low-level” properties such as color, shape, size, motion, pitch, loudness, sourness, sweetness, etc.

Even the liberal about experiential content can accept the view that perceptual experiences don’t have contents such as *not being F*, *being F or G or H*, or *being more similar to F than to G*. It is an advantage of the argument here that it does not depend on how the debate between the liberal and conservative turns out.<sup>70</sup>

## 2.2 *Argument for Immediate Justification*

---

<sup>69</sup> For a defense of liberalism, see Siegel (2006; 2010) and Bayne (2009). For a defense of conservatism, see Byrne’s contribution to Siegel & Byrne (2017).

<sup>70</sup> See Silins (2013).

Let's turn to arguing for the more difficult thesis: that the subject's belief in each case is immediately justified. I want to motivate a particular strategy for making this argument by highlighting three initial difficulties. First, I don't want to rely on any intuitions we might have about which beliefs are immediately justified and which aren't. *Immediate justification* is a term of art, and I do not want to assume that we have reliable intuitions about the structure of justification, particularly for perceptual beliefs. Second, the fact that a subject does not perform a conscious inference in forming a belief does not show that that belief is immediately justified. S's belief that p can be based on S's justified belief that q, even if S does not consciously infer the first from the second. (For example, even if I do not consciously infer that an opinion piece will have conservative view from my justified belief that it is a Fox News opinion piece, the former belief can be based on the latter.) So we cannot conclude that the mismatch subjects' beliefs are immediately justified because they are not the products of conscious inference. Finally, it would be question-begging to argue that the mismatch subjects' beliefs are immediately justified by showing how a particular view of doxastic justification predicts that they are immediately justified. In the latter parts of the paper, I'll argue that only certain views of doxastic justification predict that the mismatch beliefs are immediately justified. But we'll need some reason, independent of those views, to opt for this interpretation of the mismatch cases.

Given these difficulties, I'll opt for the following strategy. I'm concerned in this paper with evaluating how the liberal should delimit the scope of immediate perceptual justification. So I'll take my interlocutor to be someone who already accepts that we can have immediately justified beliefs in some external world propositions. That is, I take my

interlocutor to be someone who agrees that some external world beliefs can enjoy knowledge-level justification without being based on justified beliefs (1) that various skeptical alternatives are false, (2) that our experiences or perceptual beliefs are reliably formed, (3) that our experiences make likely certain propositions, (4) that it appears/seems that *p*, (where this is understood as a belief about a mental state), etc. The argument in this section, then, is directed towards those who already accept that at least some external world beliefs can enjoy knowledge-level justification without being based on any beliefs of these sorts.

However, one can be a liberal and still deny that the mismatch subjects' beliefs are immediately justified. The thought would be that we can be immediately justified in believing *some* external world propositions, but just not the ones in the mismatch cases. My strategy will be to argue against the most plausible such views. The argument will thus support 1 by casting doubt on the best ways of resisting 1 from within the liberal camp.

According to the first such view, if Anita, Ben, and Clifton are normal subjects, they had to have formed certain justified beliefs when they learned to recognize turquoise, notes that are A, C, or D, and colors that look more similar to yellow than purple.<sup>71</sup> Consider, for example, Anita's case. In learning how to recognize turquoise, we might suppose, Anita at some point saw something turquoise and justifiably believed that *that's turquoise*, or something to that effect, on the basis of some testimony. (Alternatively, she might have seen something indigo and justifiably believed on the

---

<sup>71</sup> By focusing on the "normal" case here, I'm imagining that the mismatch subjects don't learn how to recognize colors and notes by way of the machinations of some benevolent demon, who bestows the relevant knowledge or dispositions. We can take this to be a further stipulation about the cases.

basis of some testimony that *that's not turquoise*. The difference doesn't matter.)

Likewise, we might imagine that Ben, in learning to recognize notes that are A, C, or D, justifiably believed on the basis of some testimony that *that's an A* (or *that's an A, C, or D*, etc.) upon hearing a note that was in fact an A, (or was in fact an A, C, or D, etc.). Call these beliefs *perceptual learning beliefs*. I'm thinking of perceptual learning beliefs as beliefs about perceived individuals, e.g., *that's turquoise*, or *that's an A*. The general thought is that learning to recognize Fs normally requires the formation of at least some justified perceptual learning beliefs about Fs.<sup>72</sup> Thus, according to this line, each subject's belief in the mismatch cases cannot be immediately justified, for the belief depends for its justification on the subject's having been justified in holding some perceptual learning beliefs.

In response, I don't want to deny that, at least normally, learning to recognize Fs requires the formation of certain justified beliefs. However, I want to highlight two reasons for thinking that these perceptual learning beliefs do not count as mediating the justification of the mismatch subjects' beliefs in the relevant way. First, it is implausible that the mismatch subjects' beliefs—i.e., *that's not turquoise*; *the note is A, C, or D*; and *the color is more similar to yellow than to purple*—are *based on* perceptual learning beliefs. After all, the subject may no longer hold any of those beliefs; she may have forgotten about the occasions on which she learned to recognize Fs and have no beliefs about the individuals she then perceived. (Moreover, the mismatch subjects' perceptual learning beliefs needn't be about the same individuals as the target beliefs.)

---

<sup>72</sup> I don't mean there to be any suggestion that all perceptual learning involves the formation of beliefs or is otherwise cognitive.

Second, suppose one thought that the mismatch subjects' beliefs are justified because they are the output of some non-inferential competence or reliable method. It is plausible that even if the mismatch subjects acquired their competences or methods in part by forming some justified beliefs about perceived individuals, those justified beliefs do not sustain the continued possession of those competences or methods, nor their potential to output justified beliefs. The reason is the same as before. Namely, the mismatch subjects' beliefs can be justified even if we imagine that they have forgotten their perceptual learning beliefs. Thus, this first attempt to show that the mismatch subjects' beliefs are mediately justified, from within the liberal camp, doesn't seem particularly promising.

A more difficult challenge arises from the possibility that the mismatch subjects' beliefs are based on justified beliefs about how things look or sound.<sup>73</sup> For example, we might understand Anita, Ben, and Clifton as tacitly performing the following inductive inferences, respectively:

#### **Inference A**

1. The sky doesn't look turquoise.
2. If the sky doesn't look turquoise, then probably it's not turquoise.
3. Probably, the sky is not turquoise.

#### **Inference B**

1. The note sounds like either A, C, or D.

---

<sup>73</sup> For a defense of this view, see McGrath (2016, 2017).

2. If the note sounds like either A, C, or D, then probably it is either A, C, or D.
3. Probably, the note is either A, C, or D.

### **Inference C**

1. That color looks more similar to yellow than to purple.
2. If that color looks more similar to yellow than to purple, then probably it is more similar to yellow than to purple.
3. Probably, that color is more similar to yellow than to purple

For this strategy to count as a version of liberalism, we must understand the first premise of each inference as a proposition about the external world rather than the subject's mental states (McGrath 2016, 119).

I'll first consider Silins' (2011) response to this type of objection, before explaining why I don't think it is fully satisfactory for my purposes; then I'll present a more suitable response. Silins' response is to suggest that, even if one holds that the mismatch subjects' beliefs are mediated by beliefs about how things look or sound, there is still another sort of proposition that they might be immediately justified in believing, which outruns the content of their experiences. Silins suggests that we consider suppositional propositions that build information about how things look or sound into their antecedents. For example, Clifton might be immediately justified in believing that if *this* (demonstrating a perceived yellow patch) is what yellow looks like, and *that* (demonstrating a perceived purple patch) is what purple looks like, then *this* (demonstrating a perceived orange patch) looks more similar to yellow than to purple.

Clifton is immediately justified in believing this suppositional proposition; it is less plausible that his justification for believing the supposition depends on justified background beliefs about the way things look. Nevertheless, the supposition is not a content of his visual experience. The general strategy is to “pack in [to the antecedent] whatever further information one might have been relying on in the original case” (2011, 351).

I agree with Silins that these suppositional beliefs can be immediately justified. However, I think we can go further; we can argue that the mismatch subjects’ justification for their *original* beliefs is not mediated by justified beliefs about how things look or sound. Silins is primarily concerned with finding a counterexample to the Content Constraint, so any counterexample will do. On the other hand, I’m focused on evaluating which versions of liberalism make correct predictions, so it matters whether we say the mismatch subjects’ original beliefs are immediately justified. I’ll argue, *contra* the looks view, that they are.

However, so far the mismatch cases are under-described. If the mismatch subjects are competent at performing inductive inferences like the ones above, and know the premises, there seems no obstacle to holding that they have mediate propositional justification to believe the conclusions. Moreover, nothing in the description of the mismatch cases rules out that the subjects do arrive at the target beliefs by competently performing Inferences A-C, in which case their beliefs would be mediate doxastically justified. For example, Case A only holds that when Anita forms “beliefs about whether some perceived object is or isn’t turquoise, her beliefs are reliably true,” which is compatible with her forming such beliefs by competently performing Inference A. Thus,

we need to flesh out the cases to make it clear that the subjects do not form their beliefs by competently performing Inferences A-C. However, simply stipulating that the subjects do not perform Inferences A-C would beg the question.

Let's amend the cases, then, in the following way. First, suspend judgment for now on which type of method the mismatch subjects deploy in forming their beliefs. Imagine, moreover, that the mismatch subjects do count as knowing the respective premises of Inferences A-C. However, suppose that they are generally incompetent at performing inductive inferences of that sort. When they try to draw a conclusion about what properties some perceived object has on the basis of propositional knowledge about how it looks or sounds, they very often arrive at false beliefs. For example, they are prone to make inferences like the following:

1. That looks like an apple.
  2. If that looks like an apple, then probably that's an apple.
  3. So probably that is an orange.
- 
1. That note sounds like an A.
  2. If that note sounds like an A, then probably it is an A.
  3. So probably that note is a C.

Because the mismatch subjects are generally incompetent at performing these sorts of inductive inference, they cannot competently perform Inferences A-C.

We're imagining, moreover, that each subject knows the relevant premises of Inferences A-C. It is controversial whether the subjects would count, in that case, as having mediate propositional justification to believe the conclusions, since they can't competently or reliably base belief in the conclusions on belief in the premises by performing the inference. However, this doesn't matter; we can grant that subjects have mediate propositional justification to believe the conclusions. (Relatedly, we can grant that belief in the conclusion would be reasonable in light of the evidence constituted by their knowledge of the premises.) As we saw in Section 1, it is common for subjects to have both immediate and mediate propositional justification for a single belief. (I'm remaining agnostic here on whether a subject can be both mediately and immediately doxastically justified in holding a single belief.) To show that the mismatch subjects' beliefs are immediately justified, we needn't rule out that they also have mediate propositional justification to hold them, (or that belief in the conclusion would be reasonable in light of the evidence constituted by the subjects' knowledge of the premises).

At this point, one might object that there is a difficulty in fleshing out the cases in this way. Ultimately, I think the mismatch subjects have available non-inferential methods that yield immediate perceptual justification. How, then, are we to imagine them? Do they often arrive at contradictory beliefs, (e.g., *that's an apple and an orange*) by deploying both their inferential methods and their non-inferential methods in categorizing a perceived object? Do they weight the deliverances of one method more heavily? In response, I suggest that we imagine that the unreliable inductive method is deployed largely independently from other available methods. Perhaps we could imagine

that deploying the unreliable inductive method inhibits any non-inferential recognitional method available to the subjects. (Notice that, perhaps despite appearances, this maneuver does not involve stipulating that the mismatch subjects' beliefs are the outputs of non-inferential methods.)

Once we flesh out the cases in this way, it becomes clearer that the mismatch subjects' beliefs are not *doxastically* justified via Inferences A-C. After all, the subjects are incompetent at performing inductive inferences of that sort. Even when they draw a conclusion that is in fact supported by the premises (as in Inferences A-C), it is just lucky that they do so. This discussion gives us some reason to think that if the mismatch subjects are doxastically justified at all, they are immediately doxastically justified. Since it is an intuitive feature of the cases that the subjects are doxastically justified (and indeed have knowledge), we have some reason to hold that they are immediately justified. Of course, we've only cast doubt in this section on two of the ways in which one might argue that the target beliefs are mediately justified. Still, the two possibilities we have considered seem to be the most plausible ways of arguing, from within the liberal camp, that the beliefs are mediately justified. For this reason, I think the argument provides significant motivation for holding that mismatch cases are instances of immediate justification. The question now is how to explain this possibility, given that the contents of the subjects' beliefs outrun the contents of their experiences.

### **3. Trouble for Phenomenal Conservatism**

In this section, I'll consider three versions of phenomenal conservatism, and argue that each one struggles to account for the mismatch cases. Perhaps surprisingly, some of the phenomenal conservative views that struggle with mismatch cases allow beliefs with "high-level" content to be immediately perceptually justified. Consider, first, a generic version of phenomenal conservatism:

**Phenomenal Conservatism**: S has *prima facie* immediate justification for the belief that p iff it seems to S that p.<sup>74</sup>

According to PC, in the absence of defeaters, S has immediate propositional justification to believe p iff it seems to her that p.

An initial complication arises here. I've been concerned with knowledge-level immediate justification in this chapter. However, phenomenal conservatism is often taken to be a view of a weaker normative status. There is a worry, then, that phenomenal conservatives and I might be concerned with distinct normative properties. In response, I think it is plausible not only that the mismatch subjects' beliefs are knowledge-level immediately justified, and in fact constitute knowledge. They also possess other epistemic statuses that the phenomenal conservative could be targeting: The subjects' beliefs are reasonable and rational in light of their other mental states. Thus, even if we

---

<sup>74</sup> Phenomenal conservatism is often formulated as a sufficient condition along the following lines: If it seems to S that p, then S has *prima facie* justification for the belief that p (Huemer 2001; 2005; 2007). This formulation leaves open whether the principle applies to both mediate and immediate *prima facie* justification. By contrast, the formulation in the main text gives both a necessary and sufficient condition for immediate justification and is compatible with an account of mediate justification that doesn't appeal to seemings.

take “justification” in PC to be reasonableness or rationality, phenomenal conservatism struggles to handle mismatch cases, or so I’ll argue.

The content of phenomenal conservatism depends a great deal on what we take seemings to be. I’ll here treat this question briefly but see Chapter 2 for an extended discussion of seemings. Two dominant views of seemings are that they are (1) a genus of *sui generis* propositional attitudes;<sup>75</sup> and (2) dispositions to form beliefs in the absence of defeaters.<sup>76</sup> Proponents of these views agree that seemings have propositional content, and that they have a distinct phenomenal character of “felt veridicality” (Tolhurst 1998, 298). That is, in contrast to states of imagination, for example, seemings feel as though they present their propositional contents as being true.<sup>77</sup> Moreover, proponents of the dominant views agree that seemings aren’t associated with any particular source of belief; there are perceptual seemings, as when it seems to me there is a red apple before me, mnemonic seemings, as when it seems to me that I had toast for breakfast, intellectual or rational seemings, as when it seems to me that the triangle inequality is true, and introspective seemings, as when it seems to me that I have a headache.

We’re primarily interested here in perceptual seemings. I’ll argue that, on a couple ways of understanding what they are, phenomenal conservatism will have trouble explaining the mismatch cases.

---

<sup>75</sup> See, e.g., Huemer (2001; 2005; 2007) and Tolhurst (1998).

<sup>76</sup> See, e.g., Werner (2014).

<sup>77</sup> This presentational phenomenal character of seemings is plausibly distinct from the feeling of being inclined to believe the content of a seeming. One can token an intuition, for example, without feeling the least inclined to believe its content because one knows it to be false. See Chapter 2 for discussion.

First, consider versions of phenomenal conservatism on which perceptual seemings are just perceptual experiences.<sup>78</sup> On this version of phenomenal conservatism, a subject has immediate justification to believe that *p* just in case she undergoes a perceptual experience that represents that *p*. It is clear why this sort of phenomenal conservatism will struggle to explain the mismatch cases. In mismatch cases, each subject's belief that *p* is immediately justified, even though *p* outruns the content of her perceptual experience. It is true that phenomenal conservatism is a view about propositional justification. However, plausibly a subject's belief that *p* is doxastically justified only if she has propositional justification to believe that *p*. Thus, if phenomenal conservatism fails to predict that the mismatch subjects have even immediate propositional justification for their beliefs, it cannot explain how their beliefs are immediately doxastically justified.

There are versions of phenomenal conservatism, however, according to which perceptual seemings and experiences are distinct types of mental states. Let's take Tucker's (2010) *radical dogmatism* as an example. Tucker distinguishes perceptual seemings from perceptual experiences and argues that perceptual seemings, not perceptual experiences, are sufficient for *prima facie* immediate justification.<sup>79</sup> Perceptual seemings, for Tucker, are also distinct from beliefs and dispositions to believe. Rather, perceptual seemings are a *sui generis* propositional attitude, with the distinctive phenomenal character of "felt veridicality." Moreover, tokening a seeming constitutes

---

<sup>78</sup> For a version of PC that identifies perceptual seemings and experiences, see Ghijzen (2015).

<sup>79</sup> Unfortunately, Tucker (2010) calls *sensations* what I've been calling experiences, and he calls *experiences* what I would call the co-occurrence of a seeming and an experience (532). Sensations have content on Tucker's view; they are not just "raw feels." I'm just going to continue using my terminology here to avoid confusion.

recognition; the expert perceiver but not the novice might be in a position to know that *that's a Douglas fir* when she sees one because she tokens a seeming that *that's a Douglas fir* while the novice doesn't.

At first glance, it might seem as though radical dogmatism can easily handle the mismatch cases. Anita might be immediately justified in believing that the sky isn't turquoise if she bases her belief on a seeming with that content, where the seeming is distinct from her visual experience as of the sky. Ben might be immediately justified in believing that the note is either A, C, or D if he bases his belief on a seeming with that content, which is distinct from his auditory experience as of the note. Finally, Clifton might be immediately justified in believing that the swatch looks more similar to yellow than purple because that proposition seems to him to be true, and he bases the belief on the seeming. Because radical dogmatism distinguishes perceptual seemings and experiences, it can allow that the range of beliefs that can be immediately perceptually justified is not limited to the content of perceptual experience. I'll argue, however, that when we look a bit closer, radical dogmatism struggles to explain what is happening in mismatch cases.

To see why, consider the following pair of cases:

**No Experience:** Kevin looks at the table in his living room. It suddenly seems to him that there is a cat on the table, though he has no perceptual experience as of a cat there. (Moreover, he is an ordinary subject and has no reliable cat-detecting faculty that operates unconsciously.) The seeming comes from out of the blue.

**Experience:** James looks at the table in his living room, and sees his cat sleeping there. The experience causes it to seem to him that the cat is on the table.<sup>80</sup>

What does radical dogmatism predict about these cases? At first glance, it might seem that radical dogmatism predicts that both James and Kevin have justification to believe that the cat is on the table because they both token seemings with that content. However, it is open to the radical dogmatist to say that Kevin has a defeater. It is controversial exactly what the defeater would be. One possibility is to say that the defeater is constituted by Kevin's knowledge or his being in a position to know that he doesn't see a cat on the table, despite being in optimal conditions for seeing one there. There are two options for thinking about Kevin's justification in light of the defeater. On one option, possessing the defeater ensures that the seeming does not give Kevin any justification whatsoever to believe its content. On the second option, possessing the defeater is compatible with the seeming's giving him some degree of justification to believe its content, but it ensures that he lacks all-things-considered justification for that belief. (These two ways of thinking about how defeat works may drive competing intuitions about the case, i.e., an intuition that Kevin has some minimal degree of justification and an intuition that Kevin has no justification at all. Fortunately, the argument here doesn't

---

<sup>80</sup> The No Experience case is related to the cognitive penetration objection to phenomenal conservatism, developed in Markie (2005, 2006), Siegel (2012), and Bergmann (2013a). The No Experience case is not a case of cognitive penetration, but it is a case in which a perceptual seeming is formed in a defective way. Moreover, the No Experience case isn't a counterexample to radical dogmatism because the subject possesses a defeater. However, there are cognitive penetration cases in which subjects don't possess defeaters that are plausibly counterexamples to radical dogmatism. For discussion, see McGrath (2013).

depend on which of these ways of understanding defeat is best, or on which of these intuitions is correct.)

What is the relevance of the No Experience case for our purposes? The case suggests that if a subject's belief is to be justified by a perceptual seeming *and to go undefeated*, as in most ordinary cases, the seeming must be accompanied by a perceptual experience. (I'm not endorsing this view, but considering what the radical dogmatist should say, in light of the No Experience case.) Moreover, not just any perceptual experience will do. Imagine, for example, a subject whose visual experience as of white cup causes in her (for the first time) a seeming that a bell is ringing, in the absence of any auditory experience that a bell is ringing.<sup>81</sup> It is plausible that a subject who bases her belief that a bell is ringing on such a seeming would not enjoy all-things-considered justification to believe the bell is ringing. (Again, the defeater might be her knowing or being in a position to know that she doesn't hear a bell ringing, despite being in optimal conditions for hearing one.) It seems, then, that in general some relationship must hold between the perceptual experience and the seeming, in order for the seeming to give the subject undefeated justification.

This discussion suggests that, for phenomenal conservatives who distinguish perceptual experiences and perceptual seemings, it is really seeming/experience pairs that do justificatory work.<sup>82</sup> Consider, for example, the following development of this thought. (I still include the *prima facie* qualifier here, to cover defeaters the subject might possess that do not owe to the lack of an appropriate perceptual experience).

---

<sup>81</sup> I'm assuming that the visual experience has representational content and isn't merely a "raw feel," so the case is distinct from the cases discussed in Bergmann (2013a).

<sup>82</sup> Reiland (2015), Audi (2013), and Lyons (2015) discuss views of this sort without endorsing them.

**Layered PC:** S has *prima facie* immediate perceptual justification to believe p iff it seems to S that p, and S has an appropriate perceptual experience.

“Appropriate” here is just a placeholder for some relation between seemings and perceptual experiences. The idea is that not just any experience/seeming pair will be sufficient for justification: There is a need to specify some relation between experiences and seemings such that pairs satisfying this relation will yield justification. However, it will be difficult for the phenomenal conservative to specify this relation while making the correct prediction about mismatch cases.

A natural thought is that an experience is “appropriate” relative to a seeming that p just when its content includes p. However, if the argument so far is right, such a proposal cannot account for the mismatch cases. Mismatch subjects are immediately perceptually justified in believing propositions that outrun the contents of their experiences. So mismatch subjects do not token experiences that are “appropriate” in this sense. Thus, the account would predict incorrectly that mismatch subjects lack *prima facie* immediate perceptual justification for their beliefs.

A second natural thought is that an experience is “appropriate” relative to a seeming that p when its content entails p. Perhaps the content of Anita’s visual experience as of the sky entails that the sky isn’t turquoise, and so if Anita tokens a seeming with the latter content, she will count as tokening an appropriate seeming/experience pair. Even if we grant this, however, it is clear that the view is too permissive. If we allow that the content of Anita’s visual experience entails the content of

her belief—that the sky isn’t turquoise—we won’t be able to rule out that it also entails, e.g., *if p, then p*. But Anita’s visual experience as of the sky and a seeming that *if p, then p* are not an appropriate experience/seeming pair; they would not give Anita *prima facie* immediate perceptual justification to believe that *if p, then p*.<sup>83</sup> So even if this proposal makes correct predictions about the mismatch cases, it does so only by being too permissive.

Finally, on a more promising proposal discussed in Reiland (2015), drawing from Chudnoff (2013), an experience *e* is appropriate relative to a seeming that *p* just when *e* has “presentational phenomenology” with respect to *p*. For an experience to have “presentational phenomenology” with respect to *p* is for it:

(T)o putatively present you with something that is in fact a truth maker for the content of the seeming in a revelatory way. The idea is that the experience has to present the truth-maker in a way that presents its relevant look. For example, for the experience of a bishop to have a presentational phenomenology with respect to the proposition that this is a bishop it has to present you with its bishop-y look. (Reiland 2015, 526)

The thought seems to be that, for an experience to be appropriate relative to a seeming that *p*, it must present the characteristic look, sound, smell, taste, or feel of the property that the seeming attributes to the object, where on Reiland’s proposal, this is a matter of the experience’s having a characteristic sort of phenomenal character.<sup>84</sup>

---

<sup>83</sup> We could save the proposal by saying that an experience is “appropriate” relative to a seeming that *p* when its content entails *p*, but where the content of the seeming is not a logical truth. However, if we allow that Anita’s visual experience entails that *the sky isn’t turquoise*, we should also allow that it entails *the sky isn’t turquoise or q*, for any *q*. The latter proposition isn’t a logical truth, but Anita’s experience together with a seeming with that disjunctive content needn’t form an appropriate experience/seeming pair.

<sup>84</sup> We don’t have space here to discuss fully the relation between looks and phenomenal character, but for recent discussion, see Martin (2010) and Byrne (2009).

However, the resulting version of Layered PC will struggle to explain how mismatch beliefs are justified. Suppose that the mismatch subjects token the relevant seemings: It seems to Anita that the sky isn't turquoise; it seems to Ben that the note is either A, C, or D; and it seems to Clifton that swatch's color is more similar to yellow than to purple. Do the subjects' experiences have presentational phenomenology with respect to the contents of their seemings? It seems not. There is no characteristic look of things that are not turquoise; the phenomenal characters of normal visual experiences of objects that are not turquoise do not share much in common. The analogous claim seems true for the properties *being A, C, or D* and *being more similar to yellow than to purple*: The phenomenal character of experiences as of instances of those properties needn't share much in common. If there is a characteristic sound of notes that are A, C, or D, or a characteristic look of colors that are more similar to yellow to purple, they are fairly disjunctive, and so disanalogous with Reiland's bishop example. Plausibly, then, the mismatch subjects' experiences do not have presentational phenomenology with respect to the contents of their seemings. If that's right, then this version of Layered PC predicts incorrectly that the subjects' beliefs fail to be immediately justified.

#### **4. Two Strategies for Handling Mismatch**

In this section, I'll diagnose why phenomenal conservatism struggles to allow for mismatch cases. I'll use this diagnosis to highlight two strategies for handling mismatch cases from within the liberal camp and briefly argue that one of these is more promising.

One feature that the three versions of PC have in common is that they are broadly evidentialist views. They each treat the mismatch subjects' beliefs as being based on a ground, piece of evidence, or reason that is constituted or provided by a seeming, perceptual experience, or seeming/experience pair. The beliefs are immediately justified because they are based, not on a belief, but on one of these three nondoxastic states. The problem for PC arises because it has a restrictive "evidence for" relation, i.e., the relation that must hold between a belief that *p* and an experience, seeming, or experience/seeming pair for the latter to provide *prima facie* immediate justification for the former. The first proposal, for example, requires that an experience share content with the belief it *prima facie* justifies, while the third requires the experience to have "presentational phenomenology" with respect to the content of the seeming and the belief. If the argument in the first part of the paper is right, however, then these views are overly restrictive: *S*'s belief that *p* can be immediately perceptually justified even if *S*'s experience *e* doesn't represent *p*, or have presentational phenomenology with respect to *p*.

This diagnosis points the way forward. There are two general strategies for handling mismatch cases from within the liberal camp. The first general strategy—call it the *basic method strategy*—is to deny that *prima facie* immediately justified perceptual beliefs need be based on a ground, piece of evidence, or reason that is constituted or provided by a perceptual experience. This strategy is exemplified by the following family of views.

**Basic Method:** S's belief that p is *prima facie* immediately justified only if it is an output or exercise of a basic competence or a basic reliable method/process.

A competence or reliable method/process M is basic just in case exercising it in forming the belief that p does not require basing the belief that p on another belief.<sup>85</sup>

I'll make two points by way of explaining *Basic Method*. First, because views in the basic method family are broadly non-evidentialist, they needn't hold that there is any relation (semantic or otherwise) that must hold between the content of an experience, and the content of a belief, for the former to *prima facie* immediately justify the latter. The mismatch subjects' beliefs are justified not because they are based on experiences that constitute or provide evidence for them, but rather because they are outputs of basic methods or competences. For example, Anita's belief that the sky is not turquoise is an output or exercise of a basic method or competence to identify turquoise objects. Ben's belief that the note is either A, C, or D is an output or exercise of a basic method or competence to identify notes that are A, C, or D. Finally, Clifton's belief that the color is more similar to yellow than to purple is an output or exercise of a basic method or competence to identify colors that are more similar to yellow than to purple. Thus, views in the basic method family can predict that the mismatch subjects' beliefs are immediately justified.

---

<sup>85</sup> See, e.g., Lyons (2009), Goldberg (2012), Goldman (1986, 2011), and Pace (2010). We can remain neutral on how to understand the basing relation. See Korcz (2015) for an overview. By calling methods "basic," I don't mean to suggest that they are fundamental or unanalyzable.

Second, *Basic Method* does not carve out any special role for the content of experience in determining which beliefs can be immediately perceptually justified. Thus, basic method views face the familiar worry that they are too permissive, counting too many beliefs as immediately perceptually justified.<sup>86</sup> In this connection, it is important to note that *Basic Method* only gives a necessary condition on immediate justification, and so is compatible with there being further conditions that address over-permissiveness. On such a picture, beliefs that enjoy immediate justification not only must be the outputs of basic competences or reliable basic methods, but those competences/methods must also, e.g., be acquired through a meta-reliable learning process (Goldman 1986, 91-92, 115), or “have developed as a result of the interplay of learning and innate constraints” (Lyons 2009, 143), or be coherence-tested against other reliable methods (Fleisher 2019), etc.<sup>87</sup> It’s beyond the scope of this paper to evaluate these proposals in detail, but we should note that *Basic Method* is compatible with a number of strategies of addressing the over-permissiveness worry.

There is a second liberal strategy for handling mismatch cases—call it the *permissive evidentialist strategy*. On this picture, immediately perceptually justified beliefs are still taken to be based on a ground, piece of evidence, or reason that is constituted or provided by a perceptual experience. However, unlike phenomenal conservatism, the permissive evidentialist strategy deploys an evidential support relation that does not require a tight semantic connection between the content of an experience

---

<sup>86</sup> The original Truemp, Norman, and superblindsighter cases can be used to make this objection, as are modified versions of these cases discussed in Lyons (2009, 64-5). It is beyond the scope of the paper to address these familiar problems. The Basic Method strategy is compatible with a variety of solutions.

<sup>87</sup> A meta-reliable learning process is a learning process through which a subject reliably acquires reliable processes.

and the content of the belief it immediately justifies. On one possibility, S's experience *e* *prima facie* immediately justifies belief that *p* iff *e* (or its content) reliably indicates *p* (a simplified version of Alston (1988)). On another possibility, S's experience *e* *prima facie* immediately justifies belief that *p* iff *p* is the best explanation available to S for why S has *e* (a simplified version of McCain (2014)). On accounts like these, the evidential support relation allows subjects to be immediately perceptually justified in believing *p* on the basis of experience *e*, even if *p* is not among *e*'s contents. That is because *p* could constitute the best available explanation of why S has *e*, or could be reliably indicated by *e*, even if *p* is not among *e*'s contents.

Although I don't have space to evaluate these views in the fully charitable way they deserve, these two examples of the permissive evidentialist strategy are clearly not promising in their current formulation. First, it is unclear whether the explanationist proposal even gets the right verdict about the mismatch cases. It seems that we can imagine the mismatch subjects as being sophisticated neuroscientists. Thus, they might "have available" a better explanation for why they are tokening their perceptual experiences, (e.g., one involving neural firing patterns), than the contents of their beliefs. (The discussion of what makes an explanation "available" to a subject in McCain (2014, 67) doesn't clearly rule out this sort of case.) If that's right, then the proposal won't predict that mismatch beliefs are immediately justified.

The simplified version of indicator reliabilism also faces familiar difficulties. On many understandings of what reliable indication is, it is plausible that *e* could reliably indicate *p*, while failing to immediately justify belief in *p*. Consider, for example, the following case from Lyons (2009):

Suppose a mad neurosurgeon rearranges some of my neural connections while I sleep, in such a way that excessive pressure on my right big toe reliably causes a sensation of warmth on my left cheek, but for some reason, I'm prone to...[believe] that there's pressure on my right big toe.  
(64)

Moreover, imagine that nothing else causes that particular cheek sensation. Your tokening the cheek sensation reliably indicates pressure on your toe, (no matter whether one adopts a frequentist, modal frequentist, or objective chance account of reliable indication). Nevertheless, it is implausible that your cheek sensation immediately justifies belief that there is pressure on your toe. (Notice that the case is similar to the Norman/Truetemp cases, but it will be difficult for the indicator reliabilist to avail herself of any of the solutions mentioned earlier without adopting the basic method strategy.)

The upshot is that the liberal can account for mismatch cases by pursuing one of these general strategies: relax the evidential support relation that must hold between an experience and the belief it immediately justifies; or allow that beliefs are immediately justified when they are exercises or outputs of basic competences or methods. Our brief discussion in this section suggests that the latter strategy is more promising.

#### **4. Conclusions**

It is possible for a subject to be immediately perceptually justified in believing *p*, even when *p* outruns the content of her perceptual experience. Even versions of phenomenal conservatism that appeal to seemings with “high-level” content struggle to explain this possibility. Basic method accounts of immediate perceptual justification (and

to a lesser extent, permissive evidentialist accounts) are equipped to handle mismatch cases.

## References

- Alston, W. (1983). What's wrong with immediate knowledge? *Synthese* 55(1), 73–95.
- (1988). An internalist externalism. *Synthese* 74(3), 265–283.
- (1989). *Epistemic justification*. Cornell University Press.
- Alter, A. & Oppenheimer, D. (2009). Uniting the tribes of fluency to form a metacognitive nation. *Personality and Social Psychology Review* 13(3), 219–235.
- Arango-Muñoz, S. (forthcoming). Cognitive phenomenology and metacognitive feelings. *Mind and Language*.
- (2014). The nature of epistemic feelings. *Philosophical Psychology* 27(2), 193–211.
- Audi, R. (1993). *The structure of justification*. Cambridge University Press.
- (2013). Doxastic innocence: Phenomenal conservatism and grounds of justification. In C. Tucker, (ed.), *Seemings and justification: New essays on dogmatism and phenomenal conservatism*. Oxford University Press.
- Bayne, T. (2009). Perception and the reach of phenomenal content. *Philosophical Quarterly* 59(236), 385–404.
- Bayne, T. & Montague, M., eds. (2011). *Cognitive phenomenology*. Oxford University Press.
- Begg, I. M., Anas A., & Farinacci, S. (1992). Dissociation of processes in belief: Source recollection, statement familiarity, and the illusion of truth. *Journal of Experimental Psychology: General* 121(4), 446–458.
- Bergmann, M. (2013a). Externalist justification and the role of seemings. *Philosophical Studies* 166(1), 163–184.
- (2013b). Phenomenal conservatism and the dilemma for internalism. In C. Tucker, (ed.), *Seemings and justification: New essays on dogmatism and phenomenal conservatism*. Oxford University Press.
- BonJour, L. (1985). *The structure of empirical knowledge*. Harvard University Press.
- Bowden, E. M., Jung-Beeman, M., Fleck, J., & Kounios, J. (2005). New approaches for demystifying insight. *Methods*, 42, 87–99.
- Briggs, R. (2012). Interventionist counterfactuals. *Philosophical Studies* 160(1), 139–166.
- Briñol, P., Petty, R. & Tormala, Z. (2006). The malleable meaning of subjective ease. *Psychological Science* 17(3), 200–206.
- Brogaard, B. (2013a). It's not what it seems: A semantic account of 'seems' and seemings. *Inquiry* 56(2–3), 210–239.
- (2013b). Phenomenal seemings and sensible dogmatism. In C. Tucker, (ed.), *Seemings and justification: New essays on dogmatism and phenomenal conservatism*. Oxford University Press.
- Burge, T. (2003). Perceptual entitlement. *Philosophy and Phenomenological Research* 67(3), 503–548.
- Byerly, T. (2012). It seems like there aren't any seemings. *Philosophia* 40(4), 771–782.
- Byrne, A. (2005). Color and similarity. *Philosophy and Phenomenological Research* 66(3), 641–665.
- (2009). Experience and content. *Philosophical Quarterly* 59(236), 429–451.
- Byrne, A. & Siegel, S. (2017). Rich or thin? In B. Nanay, (ed.), *Current controversies in*

- philosophy of perception*. New York: Routledge.
- Carruthers, P. (2017). Are epistemic emotions metacognitive? *Philosophical Psychology* 30(1–2), 58–78.
- Chisholm, R. (1957). *Perceiving: A philosophical study*. Ithaca: Cornell University Press.
- Christensen, D. (2010). Higher-order evidence. *Philosophy and Phenomenological Research* 81(1), 185–215.
- Chudnoff, E. (2013). *Intuition*. Oxford University Press.
- Cohen, S. (2005). Why basic knowledge is easy knowledge. *Philosophy and Phenomenological Research* 70(2), 417–430.
- Cullison, A. (2010). What are seemings? *Ratio* 23(3), 260–274.
- (2013). Seemings and semantics. In C. Tucker, (ed.), *Seemings and justification: New essays on dogmatism and phenomenal conservatism*. Oxford University Press.
- Davidson, D. (1987). Knowing one's own mind. *Proceedings and addresses of the American Philosophical Association* 661, 441–458.
- Davies, M. (2000). Externalism and armchair knowledge. In P. Boghossian & C. Peacocke, (eds.), *New Essays on the A Priori*. Oxford: Oxford University Press.
- (2003). The Problem of armchair knowledge. In S. Nuccetelli, (ed.), *New Essays on Semantic Externalism and Self-Knowledge*. Oxford: Oxford University Press.
- Dokic, J. (2012). Seeds of self-knowledge: Noetic feelings and metacognition. In M. Beran, J. L. Brandl, J. Perner, & J. Proust, (eds.), *Foundations of metacognition*. Oxford: Oxford University Press.
- (2014). Feelings of (un)certainty and margins for error. *Philosophical Inquiries* 2(1), 123–144.
- Evans, J. (2009). How many dual-process theories do we need? One, two, or may? In J. Evans & K. Frankish, (eds.), *In two minds: Dual processes and beyond*. Oxford University Press.
- Fleisher, W. (2019). Method coherence and epistemic circularity. *Erkenntnis* 84(2), 455–480.
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19, 25–42.
- Ghijsen, H. (2015). Grounding perceptual dogmatism: What are the perceptual seemings? *The Southern Journal of Philosophy* 53(2), 196–215.
- Goldman, A. (1986). *Epistemology and cognition*. Harvard University Press.
- (2011). Toward a synthesis of reliabilism and evidentialism? Or: evidentialism's troubles, reliabilism's rescue package. In T. Dougherty, (ed.), *Evidentialism and its discontents*. Oxford University Press.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford University Press.
- Hawthorne, J. & Lasonen-Aarnio, M. (ms). Not so phenomenal!
- Horowitz, S. & Sliwa, P. (2015). Respecting all the evidence. *Philosophical Studies* 172(11), 2835–2858.
- Huemer, M. (2001). *Skepticism and the veil of perception*. Rowman & Littlefield.
- (2005). *Ethical intuitionism*. Palgrave Macmillan.
- (2007). Compassionate phenomenal conservatism. *Philosophy and*

- Phenomenological Research* 74(1), 30–55.
- Jackson, F. (1977). *Perception: A representative theory*. Cambridge University Press.
- Kahneman, D. & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin, & D. Kahneman, (eds.), *Heuristics and biases*. Cambridge University Press.
- Kelley, C. M. & Lindsay, D. S. (1993). Remembering mistaken for knowing: Ease of retrieval as a basis for confidence in answers to general knowledge questions. *Journal of Memory and Language* 32(1), 1–24.
- Kelly, T. (2005). The epistemic significance of disagreement. In J. Hawthorne & T. Gendler, (eds.), *Oxford Studies in Epistemology*. Oxford University Press.
- (2010). Peer disagreement and higher-order evidence. In R. Feldman & T. Warfield, (eds.), *Disagreement*. Oxford University Press.
- Knoblich, G., Ohlsson, S., Haider, H., & Rhenius, D. (1999). Constraint relaxation and chunk decomposition in insight problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25, 1534–1555.
- Korcz, K. (2015). The epistemic basing relation. In E. Zalta, (ed.), *The Stanford Encyclopedia of Philosophy*.
- Kratzer, A. (1989). An investigation of the lumps of thought. *Linguistics and Philosophy* 12(5): 607–653.
- Kung, P. (2010). On having no reason: Dogmatism and Bayesian confirmation. *Synthese*, 177(1), 1–17.
- Lewis, D. (1973). *Counterfactuals*. Harvard University Press.
- (1979). Counterfactual dependence and time’s arrow. *Noûs* 13(4), 455–476.
- Lyons, J. (2009). *Perception and basic beliefs: Zombies, modules, and the problem of the external world*. Oxford University Press.
- (2015). Seemings and justification. *Analysis* 75(1), 153–164.
- MacGregor, J. N., Ormerod, T. C., & Chronicle, E. P. (2001). Information processing and insight: A process model of performance on the nine-dot and related problems. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 27, 176–201.
- Markie, P. (2005). Epistemically appropriate perceptual belief. *Noûs* 40(1), 118–142.
- (2006). The mystery of direct perceptual justification. *Philosophical Studies* 126(3), 347–373.
- Martin, M. (2010). What’s in a look? In B. Nanay, (ed.), *Perceiving the world*. Oxford University Press.
- McCain, K. (2014). *Evidentialism and epistemic justification*. Routledge.
- McGrath, M. (2013). Phenomenal conservatism and cognitive penetration: The bad basis counterexamples. In C. Tucker, (ed.), *Seemings and Justification*. Oxford: Oxford University Press.
- (2016). Looks and perceptual justification. *Philosophy and Phenomenological Research*, 92(3), 110–133.
- (2017). Knowing what things look like. *Philosophical Review* 126(1), 1–41.
- (ms). The role of looks in human and animal knowledge: A reply to Shieber.
- Michaelian, K. & Arango-Muñoz, S. (2014). Epistemic feelings, epistemic emotions.

- Philosophical Inquiries* 2(1), 97–122.
- Moretti, L. (2015). In defense of dogmatism. *Philosophical Studies* 172(1), 261–282.
- Pace, M. (2010). Foundationally justified perceptual beliefs and the problem of the speckled hen. *Pacific Philosophical Quarterly* 91(3), 401–441.
- Papineau, D. (2006). Phenomenal and perceptual concepts. In T. Alter & S. Walter, (eds.), *Phenomenal concepts and phenomenal knowledge: New essays on consciousness and physicalism*. Oxford University Press.
- Peacocke, C. (2004). *The realm of reason*. Oxford University Press.
- Pearl, J. (2009). *Causality*. Cambridge University Press.
- Proust, J. (2008). Epistemic agency and metacognition: An externalist view. *Proceedings of the Aristotelian Society* 108(3), 241–268.
- (2013). *The philosophy of metacognition: Mental agency and self-awareness*. Oxford University Press.
- Pryor, J. (2000). The skeptic and the dogmatist. *Noûs* 34(4), 517–549.
- (ms). Uncertainty and undermining.
- Reber, R. & Schwarz, N. (1999). Effects of perceptual fluency on judgments of truth. *Consciousness and Cognition* 8(3), 338–342.
- Reder, L. (1987). Strategy selection in question answering. *Cognitive psychology* 19(1), 90–138.
- (1988). Strategic control of retrieval strategies. In G. Bower, (ed.), *The psychology of learning and motivation* 22. New York: Academic Press.
- (ed.). (1996). *Implicit memory and metacognition*. Lawrence Erlbaum Associates.
- Reder, L. & Ritter, F. E. (1992). What determines initial feeling of knowing? Familiarity with question terms, not with the answer. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 435–451.
- Reiland, I. (2015). Experience, seemings, and evidence. *Pacific Philosophical Quarterly* 96(4), 510–534.
- Rescorla, M. (2014). The causal relevance of content to computation. *Philosophy and Phenomenological Research* 88(1), 173–208.
- Ross, W. D. (1930). *The right and the good*. Oxford University Press.
- Schellenberg, S. (2008). The situation-dependency of perception. *The Journal of Philosophy* 105(2), 55–84.
- Schoenfield, M. (2015). A dilemma for calibrationism. *Philosophy and Phenomenological Research* 91(2), 425–455.
- Schwartz, B. (2002). *Tip-of-the-tongue states: Phenomenology, mechanisms and lexical retrieval*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Schwartz, B. & Metcalfe, J. (2010). Tip-of-the-tongue (TOT) states: Retrieval, behavior, and experience. *Memory and cognition*, 39, 737–749.
- Shieber, J. (2017). Looks and the immediacy of visual objectual knowledge. *Analysis* 77(4), 741–750.
- Siegel, S. (2006). Which properties are represented in perception? In T. Gendler & J. Hawthorne, (eds.), *Perceptual experience*. Oxford University Press.
- (2010). *The contents of perceptual experience*. Oxford University Press.
- (2012). Cognitive penetrability and perceptual justification. *Noûs* 46(2), 201–

- 222.
- (2017). *The rationality of perception*. Oxford University Press.
- Silins, N. (2007). Basic justification and the Moorean response to the skeptic. In T. Gendler & J. Hawthorne, (eds.), *Oxford Studies in Epistemology*, Vol. 2. Oxford University Press.
- (2011). Seeing through the veil of perception. *Mind* 120(478), 329–366.
- (2013). The significance of high-level content. *Philosophical Studies* 162(1), 13–33.
- Skene, M. (2013). Seemings and the possibility of epistemic justification. *Philosophical Studies* 163(2), 539–559.
- Sprugnoli, G., Rossi, S., Emmendorfer, A. Rossi, A., Liew S.-L., Tatti, E., di Lorenzo, G... Santarnecchi, E. (2017). Neural correlates of the eureka moment. *Intelligence*, 62, 99–118.
- Sosa, E. (1993). Proper functionalism and virtue epistemology. *Noûs* 27, 51–65.
- (1996). Rational intuition: Bealer on its nature and epistemic status. *Philosophical Studies* 81(2–3), 151–162.
- (2003). Does knowledge have foundations? In E. Sosa & L. BonJour, *Epistemic justification: Internalism vs. externalism, foundations vs. virtues*. Blackwell.
- (2007) *A virtue epistemology: Apt belief and reflective knowledge, Vol. 1*. Clarendon Press.
- (2011). *Knowing full well*. Princeton University Press.
- (2015). *Judgement and agency*. Oxford University Press.
- Thompson, V. (2009). Dual-process theories: A metacognitive perspective. In J. Evans & K. Frankish, (eds.), *In two minds: Dual processes and beyond*. Oxford University Press.
- Thompson, V., Prowse Turner, J. A., & Pennycook, G. (2011). Intuition, reason, and metacognition. *Cognitive Psychology* 63, 107–140.
- Tolhurst, W. (1998). Seemings. *American Philosophical Quarterly* 35(3), 293–302.
- Tucker, C. (2010). Why open-minded people should endorse dogmatism. *Philosophical Perspectives* 24(1), 529–545.
- Unkelbach, C. & Greifeneder, R. (2013). A general model of fluency effects in judgment and decision making. In C. Unkelbach & R. Greifeneder, (eds.), *The experience of thinking: How the fluent of mental processes influences cognition and behavior*. Psychology Press.
- Veltman, F. (2005). Making counterfactual assumptions. *Journal of Semantics* 22(2), 159–180.
- Vogel, J. (2008). Internalist responses to skepticism. In J. Greco, (ed.), *The Oxford Handbook of Skepticism*. Oxford University Press.
- Weatherston, B. (2007). The Bayesian and the dogmatist. *Proceedings of the Aristotelian Society* 107(1–2), 169–185.
- Weisberg, J. (in press). Belief in psyontology. *Philosophers' Imprint*.
- Werner, P. (2014). Seemings: Still dispositions to believe. *Synthese* 191(8), 1761–1774.
- White, R. (2006). Problems for dogmatism. *Philosophical Studies* 131(3), 525–557.
- (2009). On treating oneself and others as a thermometer. *Episteme* 6(3), 233–250.
- Woodward, J. (2003). *Making things happen: A causal theory of explanation*. Oxford

- University Press.
- (2015). Interventionism and causal exclusion. *Philosophy and Phenomenological Research* 91(2), 303–347.
- (2016). Causation and manipulability. In E. Zalta (Ed.), *Stanford encyclopedia of philosophy*.
- Woodward, J. & Hitchcock, C. (2003). Explanatory generalizations, Part I: A counterfactual account. *Noûs* 37(1), 1–24.
- Wright, C. (2002). (Anti-)Sceptics Simple and Subtle. *Philosophy and Phenomenological Research* 65(2), 330–348.
- (2004). Warrant for nothing (and foundations for free)? *Aristotelian Society Supplementary Volume* 78(1), 167–212.
- (2007). The perils of dogmatism. In S. Nuccetelli & G. Seay, (eds.), *Themes from G.E. Moore: New Essays in Epistemology and Ethics*. Oxford University Press.