UNIVERSITY OF CALIFORNIA

Santa Barbara

How To Be A Dogmatist

A dissertation submitted in partial satisfaction of the

requirements for the degree Doctor of Philosophy

in Philosophy

by

Timothy Bartholomew Butzer

Committee in charge:

Professor Kevin Falvey, Co-Chair

Professor Michael Rescorla, Co-Chair

Professor Aaron Zimmerman

Professor Anthony Bruekner In Memoriam

September 2015

The dissertation of Timothy Bartholomew Butzer is approved.

_____

Aaron Zimmerman

_____

Kevin Falvey Committee Chair

_____

Michael Rescorla Committee Chair

July 2015

ACKNOWLEDGEMENTS

Writing a dissertation is the end of a very long process (longer for me than many) and there are many people who have provided me with support and made significant contributions without which this document would not exist. I apologize in advance if I leave anyone out, as there are simply too many people who have been important and instrumental in this massive undertaking. I was introduced to philosophy as an undergraduate at the University of Colorado. I don't know how many of the faculty would remember me now, but I can honestly say that I would not be doing what I am currently doing without the fantastic experience I had in my many philosophy classes there. I particularly enjoyed those taught by Sheralee Brindell, Michael Huemer, Diane Mayer and Claudia Mills. Claudia also had the dubious distinction of serving as advisor for my honors thesis, the first major bit of writing I ever attempted. It has been years since I have seen her, but I still remember her classes, and our conversations fondly.

Upon graduating I applied widely to Ph.D. programs in philosophy, alas, to no avail. To my great fortune, however, I was accepted into a terminal Masters program at the wonderful community that is the Miami University philosophy department. It was there that I decided that studying philosophy and teaching philosophy was something that I could, and very much wanted, to do. In particular I would like to thank Asher Seidel, Emily Zakin, Gaile Pohlhaus, Peter Schuler and Elaine Miller. With the exception of Prof. Seidel, none of these scholars worked in areas that would be my focus. However, their exceptional enthusiasm and their ability to cultivate and encourage active and engaged seminars made even studying the likes of Hegel, Badiou, and Derrida enjoyable to this overly analytic mind of mine. They were tolerant of my constant requests for definitions of terms like

"diagonalization" and "The Other" and I truly enjoyed the time I spent in their courses. Just as important to my experiences at Miami University was the excellent community of graduate students I found there. Some of my fondest memories are of arguing philosophy over beers at the Circle Bar with Perry Zurn, Matthew Clark, Justin Pettibone, and Craig Howell. This was one of the friendliest and most open intellectual groups of which I have ever been a part. It was a great honor and a privilege to have studied with all of them.

My first (brief) foray into my Ph.D. studies occurred at Arizona State University. Though only there for a year, I made great friends, and grew significantly as a student of philosophy. My conversations and courses with Angel Pinillos were of particular help in this regard. I think some of my most significant philosophical development during this year came, once again, from arguments with my fellow graduate students. In particular, Parker Crutchfield, Scott Scheall, Cecilia Mun, Pete Marchetto, and Shyam Nair. I also could not have survived my year in the desert without frequent hydration and an occasional Vikings game on TV (in the corner of a Patriots bar surrounded by Patriots fans). So a special thank you to Pete and Danae Marchetto for help keeping me sane in that desert wasteland.

In 2008, I landed a position in the Ph.D. program at the University of California Santa Barbara. I want to take this opportunity to express my great appreciation for the help, guidance, support, and friendship I have received from all of the faculty, staff, and graduate students in this wonderful department. I was so fortunate in my friends and mentors, that I simply cannot, in any reasonable amount of space, do justice to the enormous amount of support I received from the department at UCSB. Thanks first to the wonderful faculty of this department. They helped me to both expand my philosophical horizons and focus in on precisely what I wanted to study. In particular I want to thank Kevin Falvey, Michael

Rescorla, Aaron Zimmerman, and Nathan Salmon. Kevin and Michael served as co-chairs of my dissertation committee and I benefited greatly from their seminars, their guidance, and their conversation. I developed many of the central ideas in this dissertation from attending their courses, and received fantastic guidance and comments throughout the composition project. Aaron became a member of my committee late, and under very difficult circumstances. I can say, without question, that this dissertation benefited greatly from his sharp eye and extensive comments. Nathan's seminars were some of the most interesting and engaging courses I attended at UCSB, and I hope to some day return to thinking about the issues we discussed (you know, once I have figured out this whole epistemology thing).

UCSB has had no shortage of brilliant and talented students during my tenure. Thanks to Alex Bundy for his friendship, commradery and constant input over the years. His epistemology reading group was one of the longest running in the department and provided many hours of entertaining and enlightening conversation over the years. I'll miss our frequent discussions over Los Arroyos burritos or Saigon curry. Thanks also to Josh May, Tim Lewis, Phil Atkins, Ian Nance and Luke Manning. All were my seniors and were very welcoming to me as a new student at UCSB They are also five of the smartest people I have ever known.. Thanks to the movie night gang: Rob McIntyre, Michael Augustin, and Kyle Dickey. Bad movies were watched, beer was consumed, and good times were had. Thanks to my roommates over the years: Matt Griffin, John Caravello, Dillon Schultz, and Corey McGrath. You might think that it would get tiresome coming *home* to a philosophical argument after reading and writing the stuff all day, but I never found it to be so. Our long meandering conversations were always a blast.

Unfortunately my last two years at UCSB were marred by tragedy. We lost an old friend and colleague as well as one of our best and brightest students. Jason Bennett was one of the most brilliant and kindest people I have ever met. He had so much promise and so much to offer the world as both a philosopher and as a person. I am honored that he thought of me as a friend. Jason's death came on the heels of another great loss. My mentor and friend Tony Brueckner's death was a massive blow to the department and to me personally. He was the most careful, intellectually honest, charitable and humble thinker I have ever known. I know that I am a better philosopher for having worked with him. Tony's intellectual charity to the arguments and individuals he encountered was matched only by his personal charity; his willingness to give his time, energy, and encouragement to his students and colleagues. He was kind, gentle, never dismissive, and one of the funniest people I have ever had the honor of knowing. Whether in conversations, seminars, reading groups, or watching the game, these times were filled with charm, humor and intellectual stimulation. He was a wonderful man, an excellent philosopher, and an amazing teacher. I miss him dearly. I dedicate this dissertation to Jason and Tony. I hope they are not too disappointed with the results.

Finally, thanks to my family. Thanks to my Mom and Dad who, though they may have had some trepidation over how long this whole thing took, gave me nothing but love and support throughout this entire endeavor. Their support allowed me the confidence to pursue a degree and ultimately a career in philosophy, despite all the challenges involved. Thanks to my brothers Eric and Scott, and to my niece and nephew Haley and Will. I love you all very much!

VITA OF TIMOTHY BUTZER
September 2015

Areas of Specialization

Epistemology, Philosophy of Mind

Areas of Competence:

Metaphysics, Philosophy of Language, Ethics

Education

University of California Santa Barbara 2008-Present

- M.A. 2011
- Ph.D. 2015

Miami University, Oxford, Ohio: M.A. 2007

University of Colorado, Boulder: B.A. Philosophy and Literature *Magna Cum Laude* 2004

Dissertation

How To Be a Dogmatist

I argue that we have immediate, defeasible warrant for our perceptual beliefs. We possess this warrant because our perceptual systems utilize assumptions and heuristic procedures—shaped by a history of successful evolutionary and developmental interactions with the environment—in order to arrive at a reliably veridical perceptual representations.

Committee: Kevin Falvey (co-chair), Michael Rescorla (co-chair), Aaron Zimmerman and Anthony Brueckner (in memoriam)

Awards and Honors

Paul Weinpahl Award for Teaching Excellence, University of California Santa Barbara 2013-2014

Graduate Students Association Excellence in Teaching Award, University of California Santa Barbara 2012-2013 (nominated but did not receive)

Steven L. Humphrey Fellowship, University of California Santa Barbara 2008-2009

Teaching Experience

Instructorships and Adjunct Faculty Positions: (University of California Santa Barbara, Arizona State University, Chandler-Gilbert Community College): Designed syllabi, exams and assignments for both lower and upper division courses, consisting of approximately 50-240 students; conducted 2-3 lectures per week; held office hours; maintained course website; supervised 1-4 teaching assistants.

PHILOSOPHY OF MIND (UCSB Spring 2015): Primarily a discussion of the developments of philosophy of mind and cognitive science in the 20th century. We discussed the merits and objections to: dualism, behaviorism, brain identity theory, functionalism, Turing Machine functionalism, and the computational theory of mind. We closed the quarter by discussing the problem of consciousness and the issues it poses for physicalist theories of the mind.

THEORY OF KNOWLEDGE (UCSB Summer 2013 and Summer 2014*)*: An introduction to central issues in epistemology, including: external world skepticism, the Gettier problem, the structure of justification, the regress of reasons, and the debate between internalists and externalists about justification.

INTRODUCTION TO ETHICS (UCSB Fall 2013): Focused on normative ethics including: divine command theory, Kant, Mill, and Aristotle. Also included select issues in applied ethics: abortion, just war theory, the permissibility of eating meat, and issues concerning poverty and charity.

INTRODUCTION TO PHILOSOPHY (UCSB Spring 2014, Spring 2013, Summer 2012): An introduction to metaphysics and epistemology. Introduced students to arguments for and against the existence of God, skepticism about the external world, the mind-body problem, and free will.

INTRODUCTION TO PHILOSOPHY (Taught simultaneously at ASU and CGCC Spring 2008): Introduced students to philosophy by focusing on issues presented in Descartes' *Meditations on First Philosophy* from both a historical and contemporary perspective. Topics included: skepticism of the external world, free will, the mind-body problem, and the existence of God.

Teaching Assistantships: (University of California Santa Barbara) Conducted two discussion sections per week for 50-70 lower and upper division students; held office hours; led review sessions; graded assignments and exams.

- Philosophy of Mind (Winter 2014)
- Introduction to Philosophy (Winter 2013)
- Philosophy of Religion (Fall 2012)
- Philosophy of Mind (Spring 2012)
- Theory of Knowledge (Winter 2012)
- Introduction to Ethics (Fall 2011)
- Theory of Knowledge (Summer 2011)
- Metaphysics (Spring 2011)
- Theory of Knowledge (Winter 2011)
- Introduction to Ethics (Fall 2010)
- Philosophy of Mind (Summer 2010)
- Critical Thinking (Spring 2010)
- Introduction to Ethics (Winter 2010)
- Introduction to Philosophy (Fall 2000)

Graduate Courses
***Indicates Audited Courses**

University of California Santa Barbara 2008-Present

Topics in Philosophy of Mind | Rescorla
Hellenistic Philosophy | Tsouna
Seminar in Epistemology | Zimmerman
Philosophy of Space and Time | Humphrey
Intermediate Symbolic Logic | Anderson
Seminar in Epistemology | Brueckner
Contemporary Philosophy | Salmon
Seminar on Hume | Holden
Meta-Ethics | Hanser
Seminar in Philosophy of Mind | Falvey
Meaning and Reference | Salmon
Seminar in Ethics | Hanser
Historical Philosophers | Zimmerman
*Seminar in Epistemology | Brueckner
*Seminar in Philosophy of Language | Salmon
*Seminar in Epistemology | Brueckner
*Topics in Theory of Knowledge | Rescorla
*Seminar in Philosophy of Language | Salmon
*Seminar in History of Philosophy | Tsouna
*Seminar in Philosophy of Language | Rescorla
*Seminar in Philosophy of Mind | Falvey
*Seminar in Philosophy of Mind | Rescorla
*Seminar in Epistemology | Zimmerman

Arizona State University 2006/2007

Philosophy of Biology | Hankinson-Nelson
Advanced Symbolic Logic | Blackson
Carnap and Quine | Creath
Modal Logic | Pinillos
Greek Philosophy | Blackson
History of Science | White

Miami University 2004-2006

19^{th} Century Philosophy | Miller
Seminar in Epistemology | Goldman
Wittgenstein | Wright
Symbolic Logic | Seidel
Agency, Identity and Social Position | Pohlhaus
Plato (Early Dialogues) | Schuller
Phenomenological Method | McKenna
Nietzsche | Miller
Metaphysics | Pohlhaus
Contemporary European Political Philosophy | Zakin

ABSTRACT

How To Be A Dogmatist

by

Timothy Bartholomew Butzer

According to the epistemological theory known as *dogmatism,* if one, for example, has a perceptual experience as of a red cube, one is immediately and defeasibly warranted in believing that there is a red cube there. The warrant is *immediate* in that it does not depend on one's other beliefs or on any introspectively accessible mental states besides the experience itself. The warrant is *defeasible* in that it is capable of being defeated by other considerations (e.g. doubts about the reliability of one's perceptual apparatus). My dissertation has three goals: (1) to defend dogmatism as a general thesis against objections, with particular focus on the arguments of Stewart Cohen, Roger White, and Crispin Wright; (2) to argue against internalist dogmatist proposals like those offered by Michael Huemer, James Pryor, and Declan Smithies; and (3) to present an original externalist theory of perceptual warrant, that I dub Competent Dogmatism, and argue that it is superior to competing accounts.

With regards to the third goal: I claim that the subject's perceptual system *itself* is normatively evaluable with regards to its ability to produce reliably veridical representations and the manner in which it does so. There are an infinite number of possible representations

of the world that are compatible with the retina being stimulated in a particular way. In order to solve this underdetermination problem, our visual systems make a set of assumptions about the environment (for example: that illumination typically comes from above, and not below). My proposed theory of perceptual warrant relies on the notion of *perceptual competence*. A subject's perceptual system is competent, if and only if: (1) the assumptions the perceptual system makes of its immediate environment are typically true of the subject's normal environment, (2) the sub-personal inferential or information processing processes conform to veridicality preserving norms, and (3) the perceptual system reliably produce veridical representations of the subject's normal environment. When one's perceptual system is competent in this sense, the representations it produces provide immediate and defeasible warrant for beliefs about one's environment. On my proposal the *normal environment* is the environment in which the assumptions and information processing processes the perceptual system uses where developed.

TABLE OF CONTENTS

Chapter I: An Introduction to Dogmatism……..1

1. Warrant and Epistemic Parsimony…….....2
2. Perceptual Warrant is Immediate and Defeasible……5
3. Clarifications and Hedges……..13
4. Dissertation Summary……19

Chapter II: Objections to Dogmatism……22

1. Introduction…….22
2. Stand-Off Defeaters……....23
3. The Revenge of Perceptually Basic Contents…….27
4. Future Evidence and Dogmatism……34
5. The Problem of Easy Knowledge……..……...……46
6. Wright's Simple Elevation Hypothesis……...56
7. The Bootstrapping Objection……..…69
8. Conclusion……..…93

Chapter III: Internalism and Dogmatism……..95

1. Metaepistemology and Dogmatism……95
2. Introspection and the Internalist Intuition……98
3. The Phenomenal Approach……..106
4. Phenomenal Conservatism……146
5. Reconsidering the Internalist Intuition……..165

Chapter IV: Competent Dogmatism……..168

1. Introduction……...168
2. Competence and Reliability……..169
3. The Underdetermiantion Problem……182
4. Solving the Underdetermination Problem…….187
5. Perceptual Competence.....196
6. Perceptual Competence and Perceptual Content……..213
7. The Contingency Problem Revisited…….220
8. The Domain Problem Again: Subjects in Skeptical Scenarios……..226
9. Conclusion…….237

References…….241

Chapter I. An Introduction to Dogmatism

In this dissertation I will mount an extended defense of dogmatism about perceptual warrant. The characteristic claim of the dogmatist is that if one has a perceptual experience as of p then one has immediate, defeasible warrant to believe p. By *immediate,* I mean that the subject's possession of the warrant does not necessarily depend on other beliefs that the subject might have. The only introspectively accessible source of the warrant, according to the dogmatist, is the experience itself. The possession of the warrant need involve no inference, even a tacit one, from other warranted beliefs that the subject possesses. According to this view one can be warranted in believing that there is a red cube in front of one without being warranted in believing that one's perceptual system is reliable, that one is not a brain in a vat, etc. In fact, one might have no attitudes towards propositions of this type whatsoever: one need not even be capable of entertaining propositions about radical skeptical scenarios or about one's own perceptual system in order to have the sort of immediate warrant that the dogmatist posits. Children, unreflective adults and even some animals will be capable of possessing this sort of warrant for their perceptual beliefs.

I will argue that these claims about perceptual warrant can be sustained in the face of the serious objections raised against them. I will also argue that a particular version of dogmatism, a view that I dub *competent dogmatism* has significant advantages over its competitors. The present chapter will serve as an articulation of the central claims of dogmatism. My discussion here will be relatively brief. The positive case for the claim that perceptual warrant is both immediate and defeasible has been ably made by many of my contemporaries and I have little to add to these discussions. My purpose in this introductory chapter is merely to establish that dogmatism is a position worthy of further investigation.

The remainder of the dissertation will provide a more rigorous argument in favor of the position by establishing that it is immune to various criticisms and by developing the view into its most plausible and defensible incarnation.

1. Warrant and Epistemic Parsimony

First, let me say something by what I mean by "warrant." This discussion is not meant to provide a definition or analysis of the term. I am dubious that any analysis of the term can be given. I will have much to say in what follows about the conditions for epistemic warrant in the domain of perception, however, here I mean simply to single out the relevant concept. I use the term "warrant" to pick out the general class of epistemic good-making features of beliefs with regards to the aim of promoting truth and avoiding falsehood. Let us call a belief epistemically warranted (or simply warranted) if it is formed in a manner that satisfies the epistemic norms that must be met in order to know something of that kind. Warrant is a necessary, though not sufficient condition for knowing a proposition. Roughly speaking, if one has a warranted, true, belief that p, (and there are no oddities of situation such as those identified by Gettier (1967)), then one knows that p. My use of the term is, therefore, consistent with that of Burge (1993, 2003) and Graham (2011).[1] Both Burge and Graham divide warrants into two categories: *entitlements* and *justifications*. Justification is a species of warrant that involves the giving and appreciation of reasons. Justifications, according to Burge (1993), must be introspectively accessible to the subject in order for the subject's

[1] Plantinga (1993) defines epistemic warrant as "whatever must be added to true belief in order to have knowledge." I am skeptical of the theoretic utility of this notion and so do not use the term in this sense.

belief to be warranted. For Burge, then, justification is a paradigmatically *internalist* form of warrant. Entitlements, on the other hand, are externalist in the sense that the subject need not be aware of warranting features of the belief in order to possess the warrant (Burge 2006). Though I will ultimately argue for an externalist dogmatist position similar to that defended by Burge, I do not mean to define alternative versions of the view into irrelevancy. There are many views that agree on the claim that perceptual warrant is both immediate and defeasible. Many of these are defended by staunch internalists some of whom, as in the case of Huemer (2006), explicitly argue that their position is vindicated by the intuitions supporting epistemic internalism. I will have much to say about the prospects of such views in Chapter III. Since I mean to take such views seriously, however, I will eschew Burge's theoretically loaded terminology in favor of the more neutral characterization of perceptual warrant as *immediate*.

I intend to adopt a stance of *epistemic parsimony* throughout the course of this dissertation. I will focus entirely on perceptual beliefs and the manner in which they are warranted. I explicitly eschew any attempt to formulate a more general account of warrant that could be applicable to other kinds of beliefs, or to other domains of knowledge. A theory such as the one I advocate could potentially be adapted to other domains, (memory seems on the face of it to be a good candidate), but I make no claims about the general adequacy of the account of warrant I discuss. There are two primary reasons for this: one is pragmatic, and the other is theoretic. Pragmatically, perceptual warrant is a complex issue that is fraught with controversy. It is difficult enough to navigate these waters without being concerned how my position may impact a general and complete epistemology that gives a full accounting of the conditions on warrant and knowledge across *all* domains. I will

regard the present dissertation as a success if it succeeds in articulating and defending a defensible account of perceptual warrant.

The theoretic reason for my parsimonious approach is that I take it to be the correct approach to epistemology generally. It is relatively common to find epistemologists who succeed in identifying a condition on warranted belief that succeeds admirably in a particular domain. Emboldened by this success, there is an understandable tendency attempt to extend this account to other domains, and in the limiting case, claim that the condition in question is a condition on warranted belief *generally*.[2] I am deeply skeptical about the prospects of such endeavors. It is certainly possible that there will be some elegant principle that, when articulated correctly, will give an adequate set of necessary and sufficient conditions for any belief to be warranted. However, I see no reason to adopt, as a constraint on theorizing about *particular* domains of knowledge and warranted belief, the assumption that such a principle is true. I find it much more likely that, certain very general features aside (for instance: that warrant has to do with "truth-conduciveness" in some very loose sense), the conditions for warranted belief could differ widely based on the features of the domain of knowledge in question. An account of warranted belief may be entirely appropriate for testimonial knowledge, but fail miserably in accounting for introspective knowledge. The epistemology of *a priori* beliefs may well present special challenges that one does not encounter with regards to scientific knowledge. In particular, though I will argue that an externalist variety of dogmatism is the most acceptable account of perceptual warrant, I do not mean this to commit me to any general claims about the structure and nature of warrant.

[2] Huemer (2001, 2006) is one such example, as is Plantinga (1993).

Despite this parsimonious approach, at various points I will confront potential problems that my chosen account of perceptual warrant that involve non-perceptual domains of knowledge. Many objections to the view I advocate attempt to show that it is incompatible with plausible principles and judgments concerning cases of non-perceptual knowledge. I do take it be a serious problem for any view of perceptual warrant if it turns out to be incompatible with plausible claims about other domains of knowledge. That is, if the view I advocate *entails* problematic positions about, for instance, knowledge of the future or being in a position to claim warrant, I regard this as a potentially serious objection. I will confront several such objections in Chapter II.

2. Perceptual Warrant is Immediate and Defeasible

i. Avoiding Over-Intellectualization

Let us begin our examination of dogmatism by considering the claim that the warrant provided by our perceptual experiences is *immediate*. A belief p is *immediately warranted* for S just in case p is warranted for S and S's warrant for p does not depend on the warrant S possesses for any other beliefs. On the face of it, this claim has a high degree of plausibility. I look at the table in front of me and see a white cylinder (in this case a coffee cup) on the table, and thereby come to warrantedly believe that there is a white cylinder there. In forming this belief, I do not enlist any other beliefs, consider what reasons I have for and against the belief, consider the general epistemic value of beliefs formed on the basis of perception, or attempt to rule out the possibility of deception or illusion. I simply look at the cylinder, and form the belief that it is there. My *perceptual experience* is a sufficient basis for my warranted belief. This relation between perceptual experiences and beliefs is, itself,

something that I very rarely reflect upon. I do not often reflect on the manner in which I form the beliefs that I have about my immediate environment; in ordinary circumstances perception and belief-formation are introspectively simultaneous. Of course, if I am asked why I believe that there is a cylinder there, I may respond by citing my perceptual experience as the basis for my belief (e.g. "Because I see it…").[3] But my ability to reflect on the manner in which my belief was formed and report on it does not entail that in my ordinary perceptions of the world I am doing something similar. Indeed, quite the opposite seems to be true. I rarely consider or reflect upon my perceptual experiences at all. Background beliefs about the epistemic status of my beliefs, my perceptual faculties and experiences, or their reliability are only considered in circumstances in which I consciously reflect on my perceptual beliefs and their causes. In ordinary cases of belief-formation based on perception my warrant for such beliefs seems irrelevant to the actual process by which my perceptual beliefs are formed.

While these considerations do not entail the immediacy of the warrant for our perceptual beliefs, they cannot be easily dismissed. One obvious rejoinder is that while I may not *consciously* base my perceptual beliefs on such background beliefs, my perceptual beliefs are still dependent on them for their warrant by virtue of *unconscious inferences* or by their relations to the justification I have for *non-occurrent beliefs* (that is beliefs that I am not currently consciously considering or reflecting on). Both strategies are problematic, and for roughly the same reasons. Small children and some animals are capable of forming beliefs on the basis of their perceptual experiences and intuitively these beliefs are

[3] In most cases I would be just as likely to simply point to the object and encourage my interlocutor to see for herself that there is indeed a red cylinder on the table.

warranted. However, such subjects often lack concepts crucial to the kinds of background beliefs usually cited as being necessary for perceptual warrant. Consider two such candidates: (1) that S's perceptual system is reliable and functioning properly, (2) that one is not deceived in some manner. Higher animals and small children will typically lack concepts such as *reliability*, *proper functioning,* and *deception* in the relevant sense. As such, it is impossible for them to possess beliefs with these contents, occurent or not. Since children and some animals are capable of possessing warranted perceptual beliefs and are incapable of entertaining propositions about the functioning or reliability of their perceptual systems, then it must not be necessary to *have* such background beliefs apply to the claim that perceptual warrant depends, in part, on unconscious inferences involving such beliefs. Once again, if it is possible for one's perceptual experiences to warrant beliefs without even the possession of such beliefs, then such warrant must not depend on unconscious inferences that involve them.[4]

These reflections concerning children, animals, and even ordinary human adult belief-formation, establish a general sort of argument against the *over-intellectualization* of the conditions for perceptual warrant. An account is over-intellectualized if it imposes constraints on warranted belief that require subjects who intuitively possess warranted beliefs of that kind to possess a degree of cognitive sophistication that they do not possess. Over-intellectualized accounts often require the possession of concepts that are beyond the

[4] One might respond that things could be different in cases in which subjects *do* possess the relevant concepts. This line of response would argue that cognitive sophistication imposes higher standards of warrant, and the possession of such concepts confers more sophisticated standards of warrant on such subjects. I respond to such a line of argument by Crispin Wright in Chapter 2, Section VI. My purpose in the present chapter is simply to present the *prima facie* case for dogmatism. I will deal with objections to the view, and elaborate the most plausible version of the view in subsequent chapters.

ken of subjects like small children and animals such as: *reliability*, *warrant, competence, perception, deception*, and *proper functioning*. If a theory requires possession of such concepts in order for a subject to be warranted in her perceptual beliefs, then the theory will be incapable of accounting for the warranted perceptual beliefs of animals, children, and perhaps even many adults. Since such subjects *do* possess warranted perceptual beliefs, such accounts will be over-intellectualized.

For instance, if the conditions a theory imposes on perceptual warrant require that a subject be able to rationally reflect on the epistemic pedigree of her perceptual experiences or to be introspectively aware of the sources of her beliefs, this will exclude many subjects who intuitively possess warranted perceptual beliefs. Rational reflection most likely entails that the subject possess complex epistemological concepts such as *reason*, *evidence, truth,* and so on. This is because rational reflection, in the relevant sense, seems to entail the ability to evaluate one's own beliefs, the rational support one has for them, and their general epistemic status. Children and small animals (and perhaps even some adults) will lack such concepts. The ability to identify the source of one's beliefs entails that one is capable of introspecting on the nature and contents of one's own mental states, such as perceptual experiences, memories, and other beliefs. But this introspective capacity seems restricted to a far smaller domain of subjects than that of warranted believers. Even if some very intelligent animals (e.g. higher mammals such as dolphins, some whales, and some great apes) possess these capabilities to some degree, it may be that dogs, cats, pigs, and horses do not possess the ability to consciously reflect on their own mental states. However, it seems clear that such subjects have rudimentary beliefs about their surroundings, and that those

beliefs are warranted. A defensible account of perceptual warrant, therefore, cannot impose constraints that entail the possession of such cognitive capabilities.

The constraint ruling out over-intellectualization of perceptual epistemology strongly favors the claim that our perceptual warrants are *immediate*. Because small children, animals, and unreflective adults form warranted beliefs without any complicated beliefs about the reliability, truth-conduciveness, or functioning of their perceptual systems, it seems as if such capabilities are not necessary for perceptual warrant. Such subjects may simply form beliefs on the basis of their perceptual experiences without basing them on any other beliefs about the world, or beliefs about the functioning or reliability of their perceptual experiences. This is not to say that such subjects cannot engage in rudimentary inferences, but rather that they need not, and often do not, do so in forming beliefs on the basis of their perceptual experiences. If one wishes to insist that perceptual warrant is not immediate in the sense articulated here, one must identify the warranted beliefs that subjects must possess in order for their experiences to warrant their perceptual beliefs. I know of no argument that establishes that *no* such grounding beliefs will be sustainable in the face of charges of over-intellectualizing the conditions for perceptual warrant. However, given the cognitive and reflective limitations of many animals and small children, it is dubious that any such belief that plausibly grounds perceptual warrant will be available to these subjects.

ii. Defeating Evidence

A subject S is defeasibly warranted in believing p just in case, S believes p, S is warranted in believing p, and it is possible that S could acquire additional evidence that would render S's belief in p unwarranted. Perceptual warrant may be defeated in many ways. One could

receive or be in possession of sufficient evidence for some proposition that is inconsistent with the contents of one's perceptual experience. A good example of such a case occurs in the Ames Room illusion (pictured below)[5].

One's visual experience of the above photograph is as of a giant child and a very small child standing about the same relative distance from the camera. However, even if you are seeing an image like this for the first time, it is unlikely that you will believe what your eyes are telling you. You are much more likely to think that the image is doctored, or that there is some other funny business going on (as indeed there is) than to believe the child on the right is as large as she appears to be. This is not due to any element of the experience itself, but rather to your extensive background knowledge about the usual relative sizes of children and rooms. Dynamic versions of the Ames Room illusion are even more striking. In such cases the subjects switch places, and in doing so one appears to grow while the other appears to shrink rapidly. Again, a normal human adult will not typically be warranted in believing that the person before them has *actually* shrunk to nearly half their previous size in the space of a few seconds in virtue of having visual experiences that represent that as occurring. This is plausibly because a normal human adult is in possession of a huge amount of evidence

[5] Image from Ames Room (2007).

that such things do not occur. Cases like this illustrate that, though having a perceptual experience as of p often suffices for one to be warranted in believing p, one could have sufficient evidence for ~*p* to override the immediate warrant one's experiences provide.

Burge (2007) argues persuasively that perceptual warrant is consistent with what he calls *brute error* (Burge 2007, 507). Brute error occurs when one's perceptual experiences are non-veridical (for instance one has an experience as of a red cube when there is not a red cube there) but the non-veridicality does not normatively implicate the perceptual system itself or the perceiving subject. Cases of perceptual illusion (such as the Ames Room) are cases of brute error. As will be explained made clear in Chapter IV, human visual systems produce representations using information that underdetermines the distal state of the world. For instance, any pattern of retinal stimulation is consistent with an infinite number of environmental states. In order to reliably represent our normal environment, our visual systems make various assumptions about the normal state of the distal environments that typically cause certain patterns of retinal stimulation. This entails that it is possible for one's visual system to be functioning properly in a manner that would, under normal circumstances, yield a veridical perceptual experience, but fail to do so due to deceptive or ambiguous features of the local environment. In such cases, according to the view I will defend, the subject is immediately warranted in her perceptual beliefs so long she is unaware of the deceptive conditions in her local environment. However, if she were to gain evidence that she was the victim of a perceptual illusion or some other deception, then this could serve to undermine or override her perceptual warrant.

This highlights another way in which one's immediate perceptual warrant can be defeated: by gaining what Pryor calls *undermining evidence* (Pryor 2013). Undermining

evidence does not present evidence against the belief in question or evidence in favor of its negation. Instead, it is evidence that the manner in which one's belief has been formed is likely to be unreliable or untrustworthy. For instance, suppose I have a perceptual experience as of a red cube in front of me in a well lit room, and I have no reason to believe that there is any funny business going on. In most ordinary cases, these conditions would suffice for me to be warranted in believing that there is a red cube here. However, suppose that my friend informs me that she has been slipping me a new hallucinogenic drug in small doses for the last several hours, and that by now I should be experiencing very realistic visual hallucinations as of red cubes. If she presents me with sufficient evidence that this is true, I may no longer be warranted in believing that there is a red cube before me. In such a case, it seems clear that if I were to continue to believe that there is a red cube in front of me given this new evidence, I would fail to be warranted in that belief. I have received evidence that has undermined my perceptual warrant for believing that there is a red cube in front of me.

A subject's immediate perceptual warrant can be undermined in many different ways. As in the case above, one could gain evidence that one's perceptual system has been compromised. One could also learn things about one's immediate environment that serve to undermine perceptual warrant. A visual experience as of a red cube could fail to provide me with perceptual warrant if came to learn that I had wandered into a technological convention for the development of holographic technology, or if I become aware of deceptive lighting conditions. My immediate perceptual warrant could be undermined *generally* if I were to become warranted in believing, via scientific investigation or in the ordinary course of experience, that my perceptual systems are systematically unreliable or incompetent. I

would no longer be warranted in simply trusting my senses if I learned that they are very likely to deceive me even in ordinary viewing conditions. In short, there are many ways for one to receive evidence that undermines or overrides the immediate warrant that one's perceptual experiences provide to believe their contents.

3. Clarifications and Hedges

The preceding arguments suffice, I think, to establish that dogmatism is worthy of further investigation.[6] Dogmatism, as I understand it here, is the claim that the perceptual warrant provided by one's perceptual experiences is both immediate and defeasible. Both claims enjoy a significant degree of initial plausibility. For a more complete positive case for the view I refer the reader to the commenters who have developed it in its various forms over the course of the past 20 years.[7] I turn now to some clarificatory remarks about the position.

First and foremost, it is important to recognize a constraint that most dogmatists place on their claims: namely that dogmatism only applies to *basic perceptual beliefs*. Though there is some disagreement about how exactly to characterize basic perceptual beliefs, the reasons for the limitation, and the general shape of the distinction between basic

[6] This very limited conclusion is all that interests me here. I will not consider other views of perceptual warrant such as: the coherentist view defended by BonJour (1985) McDowell's (1996) disjunctivist position, or Wright's (2004) neo-Wittgensteinian view, or how these views would attempt to deal with the considerations I discuss above. A comparative analysis of all competing views of perceptual warrant is a subject of such complexity and depth that it warrants a monograph. My purpose here is merely to provide some initial motivation for dogmatism in order to ground an extended discussion of the position. I take the preceding arguments to do just that without entailing that any and all competing views are false.

[7] Defenders of dogmatist positions (as I define the term) include: Burge (2003), Cullison (2010), Huemer (2001, 20006, 2007), Lyons (2009), Peacocke (2004), Plantinga (1993), Pryor (2000, 2013) and Smithies (2014).

and non-basic perceptual beliefs, are largely agreed upon. It is clear that some beliefs that are intuitively perceptual do not enjoy immediate warrant. In other words the warrant for *some* perceptual beliefs does depend on the subject possessing other warranted beliefs. Some such beliefs clearly involve some inference, either unconscious or habitual. For instance, suppose that I glance out the window at the thermometer and come to believe that it is 70 degrees outside. This belief could be warranted, and it is clearly formed (at least in part) on the basis of perception. However, in order for my belief about the outside temperature to be warranted, I must possess a number of other warranted beliefs. I must know what sort of device the thermometer is, and what it is supposed to measure. Perhaps I must have some reason to believe that the thermometer is reliable. Even if in this particular case I do not consciously consider these background beliefs in forming my belief about the outside temperature, I must possess them and be warranted in believing them in order for my temperature belief to be warranted.

Consider, also, my belief that there is a laptop computer on the table in front of me. In some sense I arrive at this belief merely by looking at the object on the table. Introspectively, I do not seem to engage in any conscious inferences, I simply see the computer and believe that there is a computer there. Thus, in some intuitive sense my belief that there is a laptop computer on the table is a perceptual belief if any belief is. However, it is unlikely that my warrant for this belief is immediate. In order to warrantedly believe that there is a computer on the table it seems I must have a host of warranted beliefs about laptop computers, their function, what they ordinarily look like, and how they can be distinguished from other objects. If I lacked warrant for any such beliefs, then my perceptual belief that there is a computer on the table would be unfounded, at best a guess as to what the object

before me is. Pryor (2000) considers the case of seeing a policeman walk down the street and forming the belief that a policeman is approaching. Such a belief, though based on perception, also seems to involve a host of background beliefs concerning policemen, how they ordinarily dress or behave, and so on.

For these reasons dogmatists typically maintain that only *basic perceptual beliefs* are immediately and defeasibly warranted. There are many ways to characterize the difference between basic and non-basic perceptual beliefs. Pryor articulates the distinction as follows:

> I will call those propositions we seem to perceive to be so, but *not* in virtue of seeming to perceive that other propositions are so, **perceptually basic** propositions, or propositions that our experiences **basically represent**. (Pryor 2000, 539)

Peacocke understands basic perceptual contents as involving kinds that are "given in the content of the experience" (Peacocke 2004, 66). Burge (2010, 546) contends that a content is perceptually basic just in case the perceptual system in question is capable of generating an experience with that content without input from higher-level cognitive processing.[8] The intuitive idea behind all of these characterizations is to, in some sense, limit the relevant class of beliefs to those whose grounding is solely dependent on the perceptual experience itself and not on any background beliefs of the subject may possess. I will characterize perceptually basic contents as follows:

[8] As we will see in Chapter II, *some* ways of understanding the distinction between basic and non-basic contents result in highly problematic consequences. These consequences do not arise for the definition given here. I will discuss these issues in more depth in considering Wright's objection to this feature of dogmatism.

> A content is **perceptually basic** for a subject just in case the content is one that the perceptual system(s) that generate the experience is capable of representing without input from higher-level cognitive processing.[9]

Precisely what contents count as perceptually basic in this sense is relative to particular subjects and is to be ultimately settled by the investigations of perceptual psychologists. However, dogmatists usually agree that the set of relevant contents is limited and involves very rudimentary concepts. Examples of contents that are usually taken to be basic on this construal include: *shadow, edge, illumination,* various color concepts like *red*, *yellow* and so on*, relative position, motion,* and various perceptible textures.[10] Perhaps some evolutionarily important contents such as *mate*, *food*, *water,* or *danger* can be represented perceptually. However, human perceptual systems are not typically thought to be capable of representing contents like *policeman* or *computer* on their own. In general most artifactual, and even many natural kinds (such as *elm tree* or *tiger*) are not likely to be the sorts of things that a perceptual system has the capacity to represent without making use of higher level conceptual resources of the subject.

As I will understand the view here, dogmatism is a view about the warrant for our *basic perceptual beliefs*, understood as beliefs formed solely on the basis of a perceptual

[9] It is likely that there is some interaction between different sensory modalities in the formation of even very basic perceptual beliefs. So, for instance, it could be the case that a particular visual experience could be generated, in part, by relying on non-visual perceptual information e.g. proprioceptive or haptic feedback. I see no principled reason that an adequate epistemological theory ought to treat such intermodal perceptual beliefs differently than those that are formed solely by use of *one* perceptual system in isolation. The above definition reflects this.

[10] See Burge (2003, 2010) for a similar definition of perceptually basic contents and arguments that contents of the sort I mention are basic in this sense.

experience and involving only the contents of the perceptual experience itself. Therefore, my preferred statement of dogmatism is as follows:

> **Dogmatism:** If one has a perceptual experience as of p, and p is a perceptually basic content, then one has immediate and defeasible warrant to believe that p.

This entails that, strictly speaking, the view I defend does not explicate the conditions for the warrant of non-basic perceptual beliefs like my belief that there is a computer on the table, or that a policeman is approaching. Indeed, this dissertation will be entirely silent concerning a huge range of beliefs that are intuitively considered to be perceptual. This limitation is necessary given the complexity and controversy of even this very limited version of the position. I hope, in future work, to construct a more complete account of perceptual warrant that extends the dogmatist thesis to the entire range of perceptual beliefs. This dissertation has the more modest aim of defending the limited version of dogmatism articulated above.

Before closing this introduction to dogmatism, I must briefly address one more issue regarding the contents of perceptual experiences. My characterization of the dogmatist position above obscures an important debate in the philosophy of perception. I have consistently stated the dogmatist position as follows: if one has a perceptual experience as of *p*, then one has immediate defeasible warrant to believe that *p*. This statement of the position entails that perceptual experiences can share contents with perceptual beliefs. However, there is disagreement in the philosophical literature as to whether perceptual experiences have *propositional* or *conceptual* contents.[11] I mean to remain neutral with regards to this debate here. While it is true that some dogmatists like Pryor (2000), Smithies

[11] There are many reasons to doubt that perceptual experiences have propositional or conceptual contents.

(2013) and Huemer (2001) assume or argue for the claim that the contents of perceptual experiences are of a kind with those in a propositional attitude such as a belief, others, notably Burge (2003), disagree. If perceptual experiences do not have conceptual contents, then this does complicate the position. A dogmatist that accepts this claim must provide some story about the transitions from perceptual to doxastic contents. It would be incumbent on such a dogmatist to reflect on the norms governing such transitions and the possible ways in which faulty or improper transitions could undermine perceptual warrant. Burge (2003) briefly discusses this transition in his discussion of the *conceptualization* of the contents of perceptual representation. I recommend his discussion to the reader.

However, the claim that one's experiences provide one with immediate and defeasible warrant for perceptual beliefs is orthogonal to this (independently interesting and important) debate about the nature of perceptual contents. Dogmatism is a claim about the structure and grounding of perceptual warrant rather than about the nature of perceptual contents. My purpose in this dissertation is to defend dogmatism as an epistemic thesis and to argue that a particular version of the view enjoys significant advantages over the to its competitors. Therefore, in the present context I intend to adopt the simplifying assumption that perceptual experiences do possess propositional or conceptual content. I will continue to characterize the position as maintaining that if one has a perceptual experience as of p, then one has immediate and defeasible warrant to believe that p. This statement of the view should not be regarded as a tacit endorsement of the claim that perceptual experiences have conceptual contents; I adopt this formulation solely because it vastly simplifies much of the discussion to follow.

4. Dissertation Summary

In the next chapter I confront prominent objections to dogmatism. In particular I argue that various arguments purporting to establish that dogmatism entails various kinds of knowledge that is *too easy* are not persuasive including: Stewart Cohen's (2002, 2005, and 2010) bootstrapping and easy knowledge arguments, as well as Roger White's (2006) argument that dogmatism entails implausible easy knowledge about the existence of future propositions. I also confront three additional arguments due to Crispin Wright (2007). The first argues that dogmatism cannot deal with a particular sort of defeater. Wright's second argument attempts to establish that dogmatisms restriction to "perceptually basic contents" is ill-founded. His third, and most fully worked out objection, is that dogmatism conflicts with plausible a plausible principle linking the conditions necessary for warranted belief and the conditions that must be met in order to *claim warrant* for a belief. I argue that each of these fails to establish its intended conclusion.

Even if one accepts that dogmatism is well motivated and can avoid major objections, a pressing question remains: what is it that explains the warrant perceptual experiences provide to our perceptual beliefs? An *internalist dogmatist*, such as Michael Huemer (2001, 2006) or Chris Tucker (2010), hold that the features of a subject's experiences that are warrant conferring must be introspectively accessible to that subject. *Externalist dogmatists*, such as Tyler Burge (2003), Christopher Peacocke (2004) and Peter Graham (2011), will allow that non-introspectively accessible features can be warrant conferring. Chapter III criticizes internalist versions of dogmatism. Internalist dogmatists typically appeal either to the *phenomenology*, or the *forcefulness* of perceptual experiences to explain their ability to immediately warrant beliefs. Against views of the first kind, I

argue that it is possible that an individual, an *uber-blindsighter*, could simultaneously have distinct experiences in different parts of her visual field with radically different qualia. Despite this qualitative difference it is possible that these distinct experiences could warrant the same perceptual beliefs for that subject. Since the phenomenology of an uber-blindsighter's experiences does not covary with the beliefs they warrant, their phenomenology cannot explain how they warrant perceptual beliefs. Views of the second kind, such as Huemer's phenomenal conservatism, claim that perceptual experiences intrinsically purport to represent the world as being a certain way. They are *assertive* in the sense that they represent their contents as actualized. This *forcefulness* is a feature of the experience that is introspectively accessible to the subject and, according to Huemer, suffices to warrant perceptual beliefs. Along with Peter Markie (2013) and Susannah Siegel (2012), I argue that a subject's perceptual experiences can be improperly influenced by strong desires, or unwarranted background beliefs of that subject. In such cases, even though the resulting experience is forceful, it is implausible that the subject can form warranted beliefs on its basis. Thus, forcefulness does not suffice to explain one's immediate perceptual warrant.

Chapter IV introduces and defends my own preferred externalist version of dogmatism. I claim that the subject's perceptual system *itself* is normatively evaluable with regards to its ability to produce reliably veridical representations and the manner in which it does so. When one's perceptual system succeeds in this regard, the representations it produces provide immediate and defeasible warrant for beliefs about one's environment. There are an infinite number of possible representations of the world that are compatible with the retina being stimulated in a particular way. In order to solve this

underdetermination problem, our visual systems make a set of assumptions about the environment (for example: that illumination typically comes from above, and not below). An experience is warranting if, (1) the assumptions one's perceptual system makes of its immediate environment are typically true of one's normal environment, (2) they have been fixed by a series of successful interactions with the normal environment and (3) they reliably produce veridical representations of one's normal environment. I argue that my view is superior to a similar account proposed by Tyler Burge (2003). Burge ties the normative evaluation of one's perceptual system to the process by which the contents of the experiences are fixed. However, the contents of one's experiences could remain unchanged while the heuristics and assumptions used by one's perceptual system change in response to changes in the environment. For instance, suppose that an organism's visual system reliably represented edges in the environment in which the contents of its experiences were fixed. If the illumination conditions of the environment change in such a way that the previous assumptions it used in detecting edges are no longer true, then its visual system may adapt by changing these assumptions. In such cases it is not the content-fixing processes that are relevant to normatively evaluating one's perceptual system, but rather the processes by which the assumptions and heuristics used by the system were developed. So long as the change in heuristics results from successful interactions with the subject's environment and allow the subject's perceptual systems to reliably produce veridical representations, these experiences will continue to provide the subject with perceptual warrant.

Chapter II. Objections to Dogmatism

1. Introduction

There are a host of objections that attempt to show that *any* dogmatist theory, regardless of the details, will fail. These arguments purport to show that the claim that our basic perceptual beliefs are warranted in the manner that the dogmatist claims lead to unsanctionable consequences. If these arguments succeed in showing that dogmatism is doomed to fail, it will hardly be worth my time to isolate, defend, and develop the best version of such a theory. I devote this chapter to showing that the objections to dogmatism presented by Crispin Wright (2007), Roger White (2006), Stewart Cohen (2002, 2005, 2010), and others do not succeed in undermining the central claims of dogmatism.

I hope the reader will excuse a somewhat lengthy survey. There are many objections that have been proposed against dogmatism in recent years, and comparatively few responses have been published on behalf of the theory. Many of these objections have not yet been subject to counter-argument in print. One major aim of this chapter is to rectify this situation by responding to several objections that seem to have (at least as far as I am aware) fallen through the cracks. I consider the objections roughly in order of ascending argumentative force: beginning with the objections that may be easily resolved, and proceeding to what I take to be the more serious problems for the view.

2. Stand-off Defeaters

Wright (2007) contends that the dogmatist account is too permissive and presents the following case to make the point:

> Judith is knowledgeably participating in a double-blind study of a new drug that she knows to have hallucinogenic properties. She also knows that half of the participants in the experiment will receive a placebo instead of the drug. The effect of the drug is to invert color vision. After the drug (or placebo) has been administered and sufficient time passes for it to take effect, Judith is shown what appears to her to be a red cube. Judith thereby forms the belief that there is a red cube in front of her (Wright 2007, 41).

In the case as described Judith is clearly not warranted in her conclusion that there is a red cube before her. For all she knows she could have received the hallucinogen instead of the placebo, in which case the cube would appear to her to be red when in fact it was green. Therefore, she ought not trust her visual system in this scenario and is not warranted in believing that the cube is red on the basis of her visual experience.

Wright maintains that dogmatism cannot accommodate this intuitive evaluation of the case. He asserts that the dogmatist must understand "reason to doubt that one's perceptual system is functioning normally" as "reason to believe that one's perceptual system is *not* functioning normally" (Wright 2007, 41). If this is correct, then dogmatism will give an incorrect evaluation of the case under discussion. It is clear that Judith doesn't have sufficient evidence to believe either that her perceptual system is working properly or to believe that her perceptual system is working improperly. Since Judith knows that there is a 50% chance that she has received the hallucinogen she ought to suspend judgment about the functioning of her visual system. Thus, according to Wright's proposed understanding of "reason to doubt," Judith has no reason to doubt that

her perceptual system is malfunctioning, and it appears to her that there is a red cube before her. It seems, therefore, that the dogmatist is committed to claiming that Judith is warranted in believing that there is a red cube before her. It is clear, however, that Judith is *not* warranted in this belief. If all of this is so, then Wright has presented a case that shows that the dogmatist theory of perceptual warrant is too permissive.

This argument relies heavily on understanding "reason to doubt" in the way that Wright does. "Doubt," Wright says, "is equivocal between *uncertainty* and *believing the negation of what is doubted* (Wright 2007, 40). Wright argues that the dogmatist cannot take the former option. His reason for this is that dogmatism "requires the status of having no reason to doubt to be consistent with having no relevant information either way" (Wright 2007, 40-41). This is correct. Dogmatists, in general, argue that a subject can be warranted in her perceptual beliefs without the subject having any attitudes or evidence whatsoever concerning the functioning of her perceptual systems. A small child may have beliefs that are warranted by her perceptual experience while lacking even the conceptual sophistication necessary to entertain propositions concerning her own perceptual systems. Let us (following Wright's terminology) characterize a position of *epistemic innocence* as follows: one is epistemically innocent with regards to a proposition just in case one has no evidence which bears for or against the proposition. It is clear that the immediate defeasible warrant that the dogmatist posits for our basic perceptual beliefs is meant to be compatible with a subject's epistemic innocence regarding propositions concerning the functioning of her perceptual system. Wright argues that this commits the dogmatist to:

> the proviso [that the subject have no reason to doubt that her visual system is functioning properly]...is met by any thinker whose evidence...provides no

> sufficient reason for a view on [the question of whether her visual system is functioning properly], and thereby mandates open-mindedness on the matter (Wright 2007, 41).

According to Wright the dogmatist must concede that unless one has *sufficient reasons* to believe that one's perceptual system is malfunctioning or compromised, then one can be warranted in believing the yields of one's perceptual system. He argues that the preceding considerations have presented the dogmatist with a dilemma: either understand "reason to doubt" in the way that Wright does, with the consequence that the view will implausibly attribute warrant in cases like Judith's, or tighten up the restrictions on what a reason to doubt could be. Wright contends that the dogmatist cannot adopt this latter tactic without being forced to reject the claim that in cases of epistemic innocence one may remain warranted in one's basic perceptual beliefs. He doubts that the dogmatist will be able to identify a relevant difference between cases of balanced evidence and epistemic innocence which will simultaneously allow them to concede that Judith is not warranted, while allowing warrant in cases of epistemic innocence (Ibid.).

So characterized, every case of epistemic innocence is a case of balanced evidence. One trivially has equal evidence for and against a proposition if one has no evidence concerning it whatsoever (Wright 2007, 41). On the other hand, all cases of balanced evidence are not cases of epistemic innocence; it is possible for one to have evidence for p, and competing evidence for ~p without the weight of the evidence favoring one over the other. Judith herself is an example of this. She certainly has *some* evidence pertaining to the functioning of her perceptual system, but she does not have sufficient evidence to believe either that her perceptual system is functioning properly or that her perceptual system is not functioning properly. The challenge for the dogmatist is

to specify a condition which successfully distinguishes cases of epistemic innocence from other cases of balanced evidence such as Judith's case, allowing the former to retain warrant while Judith does not. Understanding the condition of having a reason to doubt that one's perceptual system is functioning normally the way that Wright does precludes such a distinction. What the dogmatist requires is an alternative way of articulating this condition.

The most straightforward way to meet this challenge is to allow that the warrant for one's perceptual beliefs can be undermined by evidence that makes it unreasonable to trust one's perceptual system in a given case. Indeed, in Chapter I, this is one of the ways in which I argued our immediate perceptual warrant could be defeated. Judith certainly possesses evidence of this type: given what she currently knows there is a 50% chance that her visual system is working improperly and is unreliable with regards to discriminating colors. Since Judith has evidence that makes it unreasonable to trust her visual system, she is not warranted in believing that the cube before her is red. Allowing that immediate perceptual warrant can be undermined in this way does not entail that cases of epistemically innocent beliefs are likewise defeated. The claim is that one's dogmatic perceptual warrant will be defeated if one gains evidence that makes it unreasonable to trust one's senses in the situation in question. Positions of epistemic innocence with regards to the trustworthiness of one's perceptual system are positions in which one has no evidence whatsoever concerning the trustworthiness of one's senses. Since there is no such evidence in a case of epistemic innocence, the dogmatist is free to claim that the dogmatic warrant is preserved in such a case, while maintaining that in a case like Judith's it would be unreasonable for her to continue to trust her senses.

This is one proposal. Here are two others from the literature that also seem to dissolve Wright's dilemma. Huemer cashes out the proviso in terms of defeaters, formulating his basic principle of phenomenal conservatism as follows: "If it seems to S that p, then, in the absence of [believed] defeaters, S thereby has at least some degree of justification for believing that p" (Huemer 2007, 30). If we understand a defeater D on p as a proposition that if one believed D one could no longer rationally believe p, then this solution handles Wright's cases in much the same way as my condition above. Peacocke suggests that if one comes to believe with warrant that one's grounds are "not strong enough" to believe p, then that belief can fail to be warranted (Peacocke 2004, 62). It seems that Judith is warranted in believing that her visual experience as of a red cube is not a strong enough reason for her to believe that there is a red cube in front of her. In cases of epistemic innocence, however, since one has no evidence concerning the strength of one's grounds, Peacocke's condition is satisfied. In short, Wright's contention that the dogmatist is committed to the strong construal of "reason to doubt" is mistaken. There are many ways that the proviso in question may be cashed out, (and has been) without committing the error that Wright identifies.

3. The Revenge of Basic Perceptual Contents

Wright also criticizes a restriction on dogmatism that I discussed in Chapter I, Section 3: namely, that the only beliefs which enjoy the perceptual warrant of the kind defended by the dogmatist are beliefs with *perceptually basic contents*. Wright characterizes such contents as those that can be the proper content of a perceptual experience (Wright 2007, 45). He contends that this commits the dogmatist to cashing out what makes a content

"perceptually basic" in a way that restricts such contents to beliefs about what *seems* to the subject to be the case. If this were so it would be devastating to the dogmatist program. In general, dogmatists maintain that the limitation to basic perceptual beliefs will preserve the warrant for at least some class of beliefs about the external world and objects therein. If Wright is correct, the restriction on dogmatic warrant to perceptually basic beliefs would preclude any beliefs about the external world from being warranted in this way. So restricted, one could only possess dogmatic warrant for believing such contents as "it seems as if this is a red cube" rather than contents like "this is a red cube." If the view is constrained in the manner that Wright argues it must be, then, at best, dogmatism entails that we possess knowledge of some subset of our own internal mental states, namely the way things perceptually seem to us to be. Given that dogmatism was meant to explain the manner in which we come to be immediately warranted in our ordinary beliefs about physical objects in the external world (rather than things about the nature of our internal mental states), this is a consequence that few proponents of the theory would be willing to accept.

Wright characterizes a perceptually basic content as a content that a perceptual experience may carry as its proper content (Wright 2007, 45). This characterization is suggestive of the definition I provided in Chapter I, but unhelpful since we have been given no idea what is required in order for a content to be the "proper content" of a perceptual experience. Wright further contends that "the dogmatic cases are those where the experience's consisting in the reception of an appearance that P is a fact independent of any issue about interpretation, or the background beliefs of the subject" (Wright 2007, 45). Wright claims that in order for a content to possibly be the proper content of a

perceptual experience in this sense, it must be the case "that should such a content be false, the experience that carried it must somehow have misfired—a misperception must be involved" (Wright 2007, 46). It isn't entirely clear what Wright means by "misperception" in this context. His mention of a "misfire" invites the interpretation that a misperception involves some sort of malfunction, or other kind of interference in the workings of one's perceptual system. For instance, if neural damage causes a particular neuron to fire out of turn resulting in a perceptual experience, or if my experience is caused by the influence of a powerful hallucinogen, these would be *misperceptions*. On the other hand, if my visual system is functioning normally and properly in response to proximal retinal stimulations, there would be no misfire, or misperception even if the experience is non-veridical. For instance, if a zookeeper is successful enough in disguising a mule as a zebra, it may be the case that it is *perceptually indistinguishable* from a real zebra. There would be no misperception in this case; the mule genuinely has all the perceivable features of a zebra (it actually is colored that way, its ears are of the proper shape, it stands the correct way, etc.) If we restrict what contents count as perceptually basic in this way, then discriminating mules from zebras and vice versa is not a perceptually basic capacity.

Consider my paradigm case of a visual experience of a red cube. It is possible that a subject could have a perceptual experience as of a red cube when the cube in front of her is actually a white cube, cleverly illuminated so as to be perceptually indistinguishable from a red cube. Wright maintains that in this case, there is no misperception because, in a sense, the perceptual system functioned perfectly (Wright 2007, 48). The proximal stimulation it received is precisely the same as the proximal

stimulation it would have received if there were a genuinely red cube in front of the subject. There is no mistake on the part of the perceptual system; the problem is that the subject finds herself in an abnormal environment. In this sense, it is possible for one to have a non-veridical visual experience as of a red cube without any misperception being involved. According Wright's analysis, therefore, even the content *this is a red cube* does not count as perceptually basic. Thus, even the types of contents to which the most careful dogmatists appeal fail Wright's test. It turns out that the only contents capable of passing this test, according to Wright, are contents which concern introspectively accessible mental features of what *seems* to the subject to be the case rather than contents concerning objects in the subject's proximate external environment: It would appear that it is only by restricting perceptually basic contents...to the sphere of appearances that the difficulty can be avoided. Only then will error implicate misperception. More specifically, the needed retrenchment would seem to be that a dogmatist account of the relation between visual appearance and material world claims will be appropriate only when the latter are confined to the kind typified by: 'Here is something that looks like a red wall,' 'Here is something that looks like a hand,' and so on (Wright 2007, 47). Wright correctly notes that if the dogmatist's claims were restricted in this manner it would be devastating to the dogmatist program (Ibid).

It is fortunate for the dogmatist that Wright's criteria for determining which contents are perceptually basic is poorly conceived.[12] He claims that an experience with

[12] It is worth noting that any dogmatist position that does not recognize the restriction of warrant to perceptually basic beliefs will not even need to consider an objection along the lines that Wright urges here. There may be many reasons for doing so: Huemer (2001) and Plantinga (1993) couch their claims about perceptual warrant in broader epistemological projects according to which the existence of perceptual warrant falls out

perceptually basic content can only be non-veridical if it involves a misperception. He provides no argument for this assertion, claiming that it is a corollary of "the very idea of a perceptually basic judgment" (Wright 2007, 46). I know of no argument that supports this claim, and there are good reasons to suppose that it is incorrect.[13] Every version of dogmatism that I am familiar with concedes that the warrant that perceptual states provide is defeasible and fallible. That is, it is possible to possess perceptual warrant for a belief when the belief is false, and when it is capable of being undermined or overridden by future evidence.

Dogmatic perceptual warrant is meant to be consistent with what Burge calls *brute error*. This is a failure of veridicality which does not reflect badly on the perceptual system, its use, or its functioning (Burge 2003, 507). The claim that one may possess warrant in cases of brute error is meant to capture the following sort of situation: a subject has an experience with a content that is perceptually basic, (e.g. of a red cube) the perceptual system doing the representing is functioning normally and well (no misperception), but the resulting belief is false. This can happen when a deceptive environment is perceptually indistinguishable from its ordinary counterpart (e.g. when a

---

of their more general approaches to epistemology. Even if one is only interested in the nature of perceptual warrant, there could be other reasons for the dogmatist to eschew such a limitation: in particular, if one wanted to allow for the possibility that higher-order conceptual capacities figure in the content of our perceptual experiences. Any dogmatist who does not endorse the restriction Wright targets will be free to ignore his criticism. However, as I will argue, even dogmatist theories that are so restricted need not fear Wright's argument.

[13] To be fair, Wright cites Pryor (2001) as endorsing something like this claim. If this is Pryor's considered view then he has made an error and Wright may be regarded as presenting reasons for rejecting the limitation to perceptually basic contents *so construed*. I will present a characterization of perceptually basic contents that is not committed to Wright's claim, and thus not susceptible to his argument.

white cube is illuminated redly so as to be perceptually indistinguishable from a red cube). As Wright correctly points out, in a case such as this the perceptual system is not making any errors; the proximal stimulation the perceptual system receives is exactly the same as the stimulation it would receive if there *were* a red cube there in normal light. Burge (Ibid.) contends that in such a deceptive case if the subject is unaware of the deception, then the subject possesses dogmatic warrant for her perceptual belief that there is a red cube before her.

If, as in cases of brute error, perceptually basic contents may be perceptually represented non-veridically with no misperception involved, then for all Wright has said those contents could very well concern states of affairs in the world and not simply how things *seem* to the subject to be the case. What is required is a characterization of perceptually basic contents that produces this result. There are many candidates that meet this challenge; I will discuss one of these. Suppose that we adopt the understanding of perceptually basic contents that I advocated in Chapter I; namely, that a content is perceptually basic just in case the perceptual system(s) in question are capable of generating an experience with that content without input higher-level cognitive processing. Burge (2010, 546) suggests something along these lines, and this is consistent with the characterization given by Pryor (2001, 538-540). This characterization, does not limit perceptually basic representations to contents that may only be non-veridically represented in cases of misperceptions; in particular, such representations may also be non-veridical in cases of abnormal environments. A white cube illuminated redly may be represented by a visual system as a red cube. All that follows from this is that the visual system may generate non-veridical experiences in

certain conditions; it does not follow from this that higher-order cognitive processing or input from any other perceptual system is involved in the representation. That a perceptual experience may be non-veridical in certain cases does not entail that the perceptual system doing the representing cannot autonomously represent those contents or that some sort of interpretation or background beliefs are in play. Thus, Wright's test fails; a belief may be perceptually basic in this sense but still, be the content of a non-veridical experience in when there is no misperception.

It should be noted that what is to count as *perceptually basic* according to this proposal is largely a question to be settled by cognitive science and perceptual psychology. Both Pryor (2000) and Burge (2010) are at pains to emphasize this point. Different perceptual systems will likely differ in the contents they are capable of representing autonomously, and the issue of which contents a particular perceptual system can so represent will likely be settled by studying the actual functioning of the perceptual systems in question. As such, the characterization of perceptually basic contents I have presented here is subject to revision. However, such a characterization of perceptually basic contents along these lines is currently defensible and available to the dogmatist and Wright has not provided any argument to the contrary. The more general point is that Wright's argument will fail so long as there is *any* way of cashing out what contents are "perceptually basic" that does not have the same entailments as Wright's proposal. I have offered one proposal, but all dogmatists are not bound to it. If my proposal is unsatisfactory, then one need only settle on another way of cashing out the distinction. There is no reason to suppose that all such proposals will have the

problematic results that Wright focuses on, and no reason to suppose that Wright has produced a serious objection to dogmatism.

4. Future Evidence and Dogmatism

White (2006) argues that dogmatism when conjoined with a plausible principle about evidence for future warrant has problematic consequences. The principle in question is:

> **Meta-Warrant Principle (MWP):** If S is warranted in believing that he will be warranted in believing p, then he is already warranted in believing p. [14]

MWP has a certain degree of plausibility and it would be objectionable if dogmatism failed to be consistent with it.[15] White argues that such an inconsistency may be made manifest by considering the proposition:

> **Hands-3*** This is not a super-fake-hand.

Where:

> A super-fake-hand is a non-hand which not only looks just like a hand, but has magical powers that prevent observers from gaining independent evidence that it is not really a hand...A super-fake-hand not only eliminates all evidence of super-fake-hands, but of all other related skeptical hypotheses too...If you purchase a

[14] I re-state White's arguments and principles in my own terminology in order to maintain consistency with the rest of my discussion.

[15] I will not argue against this principle here, but I do have reservations about its truth as stated. First, it is not clear whether the justification of the antecedent is identical to the justification mentioned in the consequent. I take it to be an open question whether evidence of evidence for p is "just as good" or "just as powerful" as straightforward evidence for p.

> super-fake-hand, you can rest assured that at least no will [warrantedly] suspect that you falsely appear to have a hand. (White 2006, 537)[16]

White then considers a case in which Moore is giving a speech at a conference. His hands will be hidden behind the podium for most of the talk, but we have ample evidence that he will close his speech by waving his hands about in order to aid in making his argument. White helps himself to two stipulations: first, that we have no reason now, before the talk, to think that Moore will falsely appear to have hands, and second, that we are not now warranted in denying that Moore has super-fake-hands.

There are three collectively exhaustive possibilities for how things will go when Moore reaches the end of his speech:

> A. Moore will raise his arms at the end of his lecture and he will not appear to have hands (e.g. unbeknownst to the crowd he has recently been the victim of a serial amputator and has courageously decided to proceed with the lecture *sans* his usual hand-waving).
> B. Moore will raise his arms above the podium and appear to have hands, but we will have some reason to think that this appearance is deceptive (e.g. a jealous colleague has been passing notes around the lecture hall which claim that the possibility described in (A) has occurred but Moore has obtained very convincing rubber hands).
> C. Moore will appear to have hands and we will have no reason to think that the appearance is deceiving. (White 2006, 538)

If A is to be the case, when Moore reaches the end of his lecture, we will clearly be in as good a position as one could expect to be in for denying the truth of Hands-3*; Moore won't appear to have any hand-like things at all on the end of his arms and super-fake-hands appear to be hands. Therefore, if he does not appear to have hands, he will not have super-fake-hands. If B is to be the case, we will also be warranted in denying

[16] For the sake of argument I put aside any doubts about the possibility of super-fake-hands, though their "magical powers" are underspecified to a degree that makes such doubts reasonable.

Hands-3*. A well-made super-fake-hand would destroy any evidence that Moore lacks hands, and hence the very presence of such evidence shows that whatever Moore has on the end of his arms cannot be a super-fake-hand. Finally, if C were to be the case, then according to dogmatism the appearance will suffice to warrant us in believing that Moore has hands, and that he does not have fake hands (super or otherwise). Hence we will once again be warranted in believing Hands-3* of the things attached to Moore's wrists. Given that A-C exhaust the possibilities of the case it follows that no matter what, when Moore raises his arms at the end of his talk the dogmatist must claim that we will be warranted in believing Hands-3*.

Prior to Moore's speech, we can, as reflective agents, go through the reasoning just sketched and thus be warranted in believing that we *will* be warranted in believing Hands-3* at the end of Moore's talk. By MWP it follows that before the talk we are *already* warranted in believing Hands-3*. Either Moore will appear to have hands or he won't. If he does not appear to have hands, then we will be justified in believing that he does not have super-fake-hands because the presence of super-fake-hands entails that it appears to us that there are hands. In the case in which it appears to us that Moore has hands, either we will have defeating evidence or we won't. If we possess defeating evidence, then it follows that Moore does not have super-fake-hands because the presence of super-fake-hands entails the absence of defeating evidence. White provides the following summation:

> We began with the stipulation that while we had no particular reason to suspect that Moore falsely appeared to have hands, we were not [warranted] in believing Hands-3*. But assuming that we are [warranted] in believing dogmatism, then we are [warranted] in believing that we *will* be [warranted] in believing Hands-3*. But if this is so then we are already [warranted] in believing Hands-3*. But this contradicts our assumption (White 2007, 539-540).

White argues that dogmatism conjoined with MWP forces us to reject one or more of his assumptions about the case and invites us to conclude that it is the assumption of the truth of dogmatism that must be rejected in order to maintain consistency. A powerful argument of this form would rely on uncontroversial, unassailable, and/or obviously true premises in order to force the denial of dogmatism rather than one of White's other assumptions. If one or more of White's assumptions is suspect, then we will have reason to reject his claim that it is dogmatism, and not the suspect premise, that must be abandoned. I do not wish to contend with MWP in the present context, which leaves three assumptions we may jettison to avoid inconsistency:

(1) We have no reason to think (prior to Moore's lecture) that he will falsely appear to have hands.
(2) We are not, prior to the talk, warranted in believing that Moore does not have super-fake-hands.
(3) Dogmatism is true.

Given that the truth of (1) seems to be entailed by the stipulated features of the case, it seems we must decide between rejecting (2) and (3).

White contends that it is (3) that must be rejected, but one might think that the reasoning just rehearsed more clearly implicates (2). Recall that super-fake-hands are such that they will always prevent anyone from warrantedly believing that they are not real hands. By the very nature of super-fake-hands we can discover *a priori* that we could never warrantedly believe that there are super-fake-hands around because if they *were* around they would destroy any reasons we could have for thinking that they were there. Knowing this, we can conclude that if we *do* acquire any purported evidence in favor of the proposition that there are super-fake-hands around, that it must be deceptive.

Furthermore, we also know that if Moore has super-fake-hands, they will provide us with every reason that normal hands would provide for believing that they are hands. The reasoning just rehearsed does not depend on or require the assumption of dogmatism. It requires only a reflection on the nature of super-fake-hands to conclude that one could never be justified in believing that Moore will have them.

Of course, the impossibility of possessing positive evidence for a proposition does not entail that we have evidence *against* that proposition. It could be the case that the only rational attitude we can adopt towards Hands-3* is a suspension of judgment. This is consistent with White's claim that, prior to the talk, we are not warranted in believing Hands-3*. However, consider a complicating feature of White's case that I have thus far ignored: in order for the subjects in his case to reason as he suggests they can, they must warrantedly believe that dogmatism is true.[17] Suppose, for the sake of argument, that I have presented arguments that make my belief in the position at least *prima facie* warranted. I can then reason as follows: before Moore begins his talk either Moore will appear to have hands at the end of his talk or he won't. If he does not appear to have hands at the end of his talk, we will then be warranted in believing that he does not have super-fake-hands because the Moore's having super-fake-hands at that time entails that it will appear to those that who view the ends of Moore's arms that he has hands. In the

[17] White is well aware of this and incorporates the claim into his argument by claiming that the assumption that the subjects were warranted in believing the truth of dogmatism must have been false. As he presents his argument, this shows that dogmatism must be false because it is an *a priori* knowable thesis, and as such if it were true, those of us who believe it would be warranted in believing in doing so. I have ignored this version of the argument because I think it does White no favors. It is at best contentious that, in general, we are all warranted in believing all *a priori* knowable theses that we believe. As I have indicated in my presentation, however, I do not think White needs this contentious claim in order to make his case against the dogmatist.

case in which it appears to us at the end of the talk that Moore has hands, either we will have defeating evidence of sufficient power to undermine our perceptual warrant or we will not. If we possess such defeating evidence, then it follows that Moore does not have super-fake-hands because the presence of super-fake-hands entails the absence of any defeating evidence whatsoever. Finally, if Moore appears to have hands at the end of his talk and there is no reason for me to doubt that he does, then we will be warranted in believing that he has hands and hence will be warranted in believing that he does not have super-fake-hands. Therefore, by MWP I am now, (before the talk) warranted in believing that Moore does not have super-fake-hands. Hence, (2) is false, and White's argument against dogmatism fails. White provides no argument in support of (2), simply taking it for granted as a premise of his argument; I suggest that a dogmatist ought to reject this assumption.

White argues that there is a problem for any dogmatist that responds to his case in this way. Let a skeptical scenario have Intermediate Epistemic Status (IE-Status) for me at a given time if at that time I am neither warranted in believing that it obtains nor am I warranted in believing that it does not obtain. White acknowledges that the dogmatist could argue, as I have, that given the nature of super-fake-hands, Hands-3* does not have IE-Status for me before the talk. That is, I can warrantedly conclude that I could not fail to have warrant for Hands-3*, and hence no actual perception is required in order to be warranted in believing the proposition. White goes on to argue that this response cannot save the dogmatist, because an argument similar to the one presented to show that I am warranted in believing Hands-3* prior to any hands-experience may *also* be marshaled to

show that I have warrant to believe that Moore does not have fake-hands prior to any investigation into the matter.

Consider the more mundane proposition:

Hands-3: That is not a fake-hand.

White constructs a new case in which I have ample evidence for believing that I will only have a very limited opportunity to uncover evidence of hand-illusions at the end of Moore's speech (Moore will only raise his arms briefly and leave quickly afterwards with no opportunity for questions or investigation). Thus I know that (B) will not be the case, leaving the only other possibilities as (A) and (C); either Moore will not appear to have hands at all or he will appear to have hands and I will have no cause to doubt this. The dogmatist will claim that both of these possibilities will result in my having warrant for believing Hands-3. In the former case, I will not have an experience as of anything that looks like a hand. Fake-hands resemble hands, and so I will have *prima facie* warrant that Moore doesn't have fake-hands. In the latter, the undefeated visual appearance as of hand will provide *prima facie* warrant to believe Hands-3. Therefore, I can conclude that whether (A) or (C) holds, I will be warranted in believing Hands-3 at the end of Moore's talk. By MWP I am now, before the talk, also warranted in believing Hands-3. Since I am warranted in believing that Hands-3, the dogmatist must say that Hands-3 does not have IE-status for me before the talk.

White draws a very general conclusion from the reasoning just rehearsed: that dogmatism is committed to the consequence that *no* skeptical alternative to Moore's having hands can have IE-status before Moore's talk. White's argument is a bit opaque, so I quote it at length:

> Now if we reflect on this case, it seems clear that the availability of this justification for Hands-3 could not really depend on our justification for expecting not to obtain any evidence that Moore falsely appears to have a hand. Our reason for expecting no such evidence to show up was just that there was no opportunity for us to obtain it, whether such evidence is available or not. This does not amount to evidence that Moore is not fake-handed. So if we grant Hands-3 in the case as described, we must admit that we would also have it without justifiably ruling out our obtaining evidence of misleading hand-appearances. The upshot is that *no* skeptical alternative to [Moore's having hands] can have IE-status for us. If we have no reason to suspect that any skeptical alternative to [Moore's having hands)] is true, then independently of [an experience of hands], we are justified in denying all skeptical alternatives. (Ibid. 541-542)

If White's argument is cogent, it does present a problem for dogmatism. It is not incumbent on the skeptic to provide evidence that her chosen scenario actually obtains. Arguments for skepticism about our knowledge of the external world typically rely on the plausibility of (e.g.) the claim that we cannot rule out the possibility that we are BIV's (brains in vats stimulated by a supercomputer to have the perceptual experiences that we do) by *a priori* reasoning similar to the kind given above. In other words, the scenarios that are typically invoked in arguments against our knowledge of the external world are generally thought to have IE-status for us. However, if all that is required for a skeptical alternative to lack IE-status is that we lack reasons to suspect that they are true (or likely, or in some sense "relevant" alternatives to what we are inclined to believe, etc), then no standard skeptical alternative will have IE-status because we are never given any reason to suppose that they actually obtain. Denying that standard skeptical scenarios have IE-status for us, while not strictly inconsistent with dogmatism, does present a theoretic cost for those who want to solve rather than dissolve standard skeptical arguments.

White's argument that the dogmatist must deny Hands-3's IE-status depends on his contention that the dogmatist must allow, in the case as described, that our warrant for

Hands-3 does not depend "on our expecting not to obtain any evidence that Moore falsely appears to have hands" (Ibid.). However, it is clear that the evidence we have before Moore's talk that Moore will not misleadingly appear to have hands *does* play such a role in the case. Recall that the dogmatist characteristically maintains that the warrant for one's basic perceptual beliefs is both fallible, and defeasible. That is having this warrant is compatible both with one's belief being false, and one's coming to possess evidence against it at some future time. But in constructing his revised case White has effectively eliminated the second possibility; I have evidence that I will not come to possess defeating evidence for my perceptual experiences. This is relevant to whether the dogmatist will claim that one is warranted in believing the propositions in question prior to having any experience.

Suppose, that I am sitting in the lecture hall before the talk and I have no evidence about what the ends of Moore's arms will look like at the end of the talk, nor about the conditions under which I will be able to view them, nor about any additional evidence I might acquire by whispered gossips during the talk. I don't know how long Moore will hold his hands aloft at the end of the talk, nor do I know whether I will see his unobstructed arms close up, or anything of the kind. In short, all I know is that Moore will wave his arms about at the end of his talk (I have seen this performance before). Am I, according to the dogmatist, right now warranted in believing that Moore does not have fake hands? Clearly not; Moore may have very obvious fake hands (cyborg hands), his prosthetic hand may fly off as he waves it about, some jealous colleague may spread the tale that Moore's hands are fake despite their realistic appearance, I may approach Moore to shake hands in congratulation of his stunning demonstration after the talk has ended

thereby revealing that the hands are convincing fakes, and so on. Any number of things might happen that would undermine any warrant that my visual experiences may initially furnish. Not being able to rule out any of these occurrences, I do not now have warrant to believe that Moore does not have fake hands. But if, as in White's case, I am warranted ahead of time in believing that none of these things will happen and I am also warranted in believing that if Moore appears to have hands, the appearance will be convincing, I certainly can be warranted now in believing that Moore does not have fake-hands. The dogmatist should maintain that in the cases White describes the evidence precluding future counter-evidence is central to whether or not MWP will generate the results that White contends it does.

Therefore, the dogmatist is committed to no general claims about all skeptical scenarios, and can consistently maintain that some of them have IE-status for us. One can certainly construct a case without IE-status by stipulating the impossibility of ever gaining evidence that one is deceived. So, for instance, one could imagine an evil demon who necessarily and with perfect success provides one with deceptive experiences. For such a scenario one can, indeed, reason as White suggests. Since I know that I will never receive any evidence that such a scenario obtains and (on the assumption of dogmatism) I also know that I will receive evidence to the contrary (namely that I have a hand and not a demon hand), I can reasonably believe, *right now* that such a scenario does not obtain. Such a demon-scenario does not, on the dogmatist proposal have IE-status for us. However, these cases are similar to the case of super-fake-hands in that their construction entails that one could never be warranted in believing that they obtain. The dogmatist's

counterargument that such a result is not problematic given their theoretic commitments seems equally strong.

One only generates a skeptical scenario without IE-status by stipulating the impossibility of ever obtaining evidence that the scenario obtains. But standard skeptical scenarios are consistent with a subject gaining evidence that they obtain. For instance, suppose that an outside benefactor has taken control of the computer generating the BIV's experiences with the goal of apprising the BIV of its plight. The benefactor introduces a ticker bar along the bottom of the BIV's visual field with the repeated message: "You are a brain in a vat. Believe nothing that you experience." If the BIV remains unconvinced, the ticker could read "in five seconds you will have experiences as of being in the desert for one minute. Then you will have experiences as of being back in your living room." If the BIV subsequently had just those experiences (and the process was repeated several times) this would certainly constitute some evidence that she was a BIV. O.K. Bouwsma (1949) imagines a demon scenario in which the demon, bored with constant unacknowledged deception of his subject Tom, attempts to apprise Tom of his deceived status by engaging him in conversation. Though Bouwsma's demon is unsuccessful in convincing Tom, surely such an exchange provides Tom with *some* evidence that the demon scenario does not obtain. Neo of the Matrix receives evidence that he has lived his entire life in an envatted state by being forcibly removed from it. One cannot rule out, in advance, that one will receive evidence that one has been massively deceived. Thus, it is compatible with dogmatism to maintain that any such

skeptical scenario has IE-status for us right now. White's conclusion that dogmatism is committed to claiming that *no* skeptical alternative has IE-status is incorrect.[18]

Of course, one could attempt to argue that our past experiences do give us good reason to suppose that we will not suddenly experience a ticker at the bottom of our visual field, or suddenly wake up to a post apocalyptic wasteland; just as our past experiences could give us good reason to suppose that Moore will not wave around a cybernetic hand at the end of his lecture. A dogmatist of a certain persuasion could seek to argue against external world skepticism by claiming that, in the absence of such defeating evidence and our extensive past experiences as of normal physical objects, we are *right now* warranted in believing that no such scenario obtains. I will remain neutral here as to the prospects of such anti-skeptical strategies. For present purposes, two points about such arguments are salient. First, they rely on *past experiences* that a subject has had. The problem White tried to present for dogmatism seemed to revolve around the claim that the dogmatist is committed to the claim that *prior to any experience* of the world, no skeptical scenario has IE-status for us. The preceding points illustrates that dogmatism entails no such conclusion.

Second, an anti-skeptical argument of the envisioned variety would attempt to establish that our past experiences allow us to gain knowledge or warranted belief that we will receive no evidence that we are massively deceived. I expect that the battle lines concerning this claim will precisely mirror those concerning the claim that our past experiences give us evidence that there are external objects. Those inclined to emphasize the strength of skeptical arguments will be similarly inclined to argue that our past

[18] Thanks to Aaron Zimmerman for helpful comments and questions on these points as presented in a previous draft of this chapter.

experiences give us no such evidence, while those of a Moorean (or perhaps a Russellian) persuasion will insist that they provide a suitable inductive or abductive base for the anti-skeptical conclusion. Perhaps reframing the debate in these terms will provide some additional support for a dogmatist anti-skeptical argument; this is an avenue of research worth pursuing. However, the prospects of such a strategy are unclear. Furthermore, if such an anti-skeptical argument is made available by the dogmatist's claims, it is far from obvious that this is objectionable. Certainly, such an argument will not have the trivial or dismissive character of the arguments against the IE-status of skeptical scenarios that White considers objectionable.[19]

5. The Problem of Easy Knowledge

Cohen (2002, 2005, 2010) argues that the conjunction of the claim that perceptual warrant is immediate and a plausible closure principle entails knowledge (or for our purposes, warranted belief) that is *too easy*. Cohen presents a principle similar to the following plausible claim (once again I alter the principle in order to keep the discussion focused on *warrant* rather than *knowledge*):

> **Closure of Warrant (CW):** If S is warranted in believing p, p entails q, and S competently deduces q from p, then S is warranted in believing q.

[19] It should also be noted that if one adopts an externalist dogmatist position of the kind I advocate in Chapter IV, all of the typical sources of dissatisfaction with externalist responses generally will serve to undermine an anti-skeptical argument along the lines I described here. For instance, if one's perceptual experiences only provide warrant if they are produced by a reliable perceptual system, the skeptic will argue that one may not know whether or not those conditions are met, or that it is question-begging to claim that they are. Regardless of the strength of such responses on behalf of the skeptic, I do not foresee that recasting the debate in the terms White imagines change the fundamental issues that separate epistemic externalists and skeptics.

His example focuses on the conditions that must be met in order to know that local deceptive conditions do not obtain. He imagines shopping for a table with his son, who wants a red table for his room. Cohen locates a table in the store that visually appears to him to be red. On the basis of this appearance he warrantedly concludes that it *is* red. Furthermore, since the table's being red entails that the table is not white and illuminated redly, CW allows Cohen to warrantedly conclude that the table is not white and illuminated redly. Cohen notes that, according to views like dogmatism which maintain that a perceptual experience as of p can warrant a belief in p without background knowledge of one's reliability or ruling out possibly deceptive conditions in one's environment, the warrant that one has for the unlikelihood of such a deception is irrelevant to the warrant one possesses for the belief. All that is required to generate the problem is that the subject possesses warrant for his perceptual belief, and that CW holds.

Cohen calls the result that he can be warranted in believing (indeed know) that the table is not white but illuminated redly in this way "a strange result" (Cohen 2002, 313). He demonstrates this strangeness by considering a dialog with his son:

> We go in the store and I say, "That table is red. I'll buy it for you." Having inherited his father's obsessive personality he worries, "Daddy, what if it's white with red lights shining on it?" I reply, "Don't worry— you see, it looks red, so it is red, so it's not white but illuminated by red lights." Surely he should not be satisfied with this response" (Ibid. 314).

Since Cohen believes that one cannot come to be warranted in believing that the table is not white and illuminated redly in this manner, he claims that views which claim otherwise fail. In particular, he claims that positions that deny what he calls the KR principle are implicated by this example, (again, I rephrase his principle to stay consistent to my focus on warranted belief):

> **WR:** A potential source of warranted beliefs K can yield warranted beliefs for S, only if S knows (or is warranted in believing) K is reliable.

Indeed, if his conclusions about his case are correct, then it will present a problem for any WR denying view for the reasons that he mentions.[20] Dogmatism denies WR by claiming that in order to be warranted in a basic perceptual belief in *p* one need only have a perceptual experience with the content *p*, and no positive reasons to doubt that belief. The dogmatist explicitly denies that the warrant for one's basic perceptual beliefs depend on the knowledge that one's perceptual systems are reliable. Therefore, if Cohen's easy knowledge problem is genuine, it will constitute an objection to dogmatism.

The main support for Cohen's claim that this case presents a problem for dogmatism emerges from a reflection on the dialog he imagines having with his son. There is a well-worn objection to using dialogs of this kind to show something about the conditions necessary for warranted belief: namely, that a subject's lacking a dialectically effective argument for a belief does not entail the subject lacks a warranted belief. There is a clear distinction between one's being warranted in believing p and one being capable

[20] As Silins (2008) points out, dogmatism is not, strictly speaking, incompatible with WR as stated. If one holds that we always have a sort of default warrant to believe that our perceptual systems are reliable, one may also hold that the warrant we have for our perceptual beliefs is immediate so long as one maintains that the perceptual warrant is not based or grounded on the default warrant concerning the reliability of one's perceptual system. The conditional claim of WR would be true, because the antecedent would always be true. It is unclear what the motivations for such a view would be, and it is clear that Cohen wishes WR to be read stronger than simply a material conditional. In other words, the views he argues against will deny that the warrant for our perceptual beliefs *depends epistemically,* on the prior warrant one has for believing in the reliability of one's senses. I ignore this complication in what follows. Thanks to Kevin Falvey for bringing this point to my attention.

of justifying or defending one's belief in a dialectically effect manner to an interlocutor.[21] As such we should not conclude that Cohen's inability to answer his son in a non-question-begging manner entails anything about whether Cohen is warranted in believing the table to be red, or the conditions he must meet in order to be so warranted. The intuitively problematic nature of Cohen's response to his son's challenge is explained by the apparent dialectical failings of Cohen's response, but this does not suffice to impugn the dogmatist's commitment to Cohen's belief being warranted.

Cohen acknowledges this point in response to a similar argument presented by Markie (2005), but does not think that it can completely explain our intuitive response to his case. He argues that his responses to his son do not indicate a mere dialectical failure in virtue of the fact that we (who are not ourselves part of the conversation) ought to find his reasoning in the case suspect. To make this point he removes his son from the case and considers a first-person version of the dialectic.

> Suppose *I* am interested in buying a red table and am anxious to avoid buying a white table that looks red (in the showroom) because it is illuminated by red lights. In order to convince myself that a particular table is not deceptively illuminated, I appeal to (the following) reasoning: "Hmmm, the table looks red, so it is red, so it is not white with red lights shining on it." Suppose further that as a result of going through this reasoning, I become convinced that the table is not white with red lights shining on it. Presumably, the intuition remains that there is something wrong with my having reasoned this way (Cohen 2005, 419).

In this first-person case, Cohen has doubts about the color of the table and is determined to "convince himself" that the table is not white and deceptively illuminated by reasoning as he does. However, in light of his doubts, it seems clear that the reasoning he marshals is not a proper basis for his belief that the table is not white and illuminated redly. In

[21] For clear arguments and examples establishing this point, see Alston (1980).

order allay those doubts he must carry out an investigation of the lighting conditions in the store in order to ascertain that he is not being deceived. This claim, however, entails nothing about what Cohen could warrantedly believe in a situation in which he did not doubt the color of the table or suspect that he may be subject of some deceptive illumination. The dogmatist need only claim that *in cases like that* Cohen warrantedly believes that the table is not white and illuminated redly. Cohen's first-person version of the case does not bear on this claim. To illustrate this point, consider a case in which we remove the doubt implicit in Cohen's presentation of the case. Cohen is shopping for a table. He looks at a table that appears to be red, and forms the belief that it is red. Cohen does not consider the lighting conditions in the store at all, but undertakes to derive some trivial logical consequences of his belief that the table is red, reasoning as follows:

(1) The table is red.
(2) Therefore, the table is not white.
(3) Therefore the table is not white and illuminated bluely.
(4) Therefore, the table is not white and on the moon.
(5) Therefore, the table is not white and made of gorgonzola.
(6) Therefore, the table is not white and directly in front of me.
(7) Therefore, the table is not white and illuminated redly.

Is there something more problematic about (7) than (2)-(6) in this chain of reasoning? It is hard to see what this could be. (7) has been arrived at by exactly parallel reasoning as (1)-(6). Moreover, it seems that the warrant for (1)-(6) can be undermined in similar ways as (7). Suppose Cohen worried about whether or not a complex series of mirrors was currently causing him to mistake the location of the table. This doubt would make the previously acceptable inference to (6) problematic. Suppose that Cohen had previously heard about an artist making realistic gorgonzola statues of common pieces of furniture. This could cause him to doubt or worry about the inference to (5). When we

evaluate the case from our idealized perspective, we take these doubts into account as part of Cohen's total epistemic state. Given these doubts, it would be irrational for Cohen to make the otherwise unproblematic inference. We can reach this conclusion even if Cohen's doubts are themselves irrational. He perhaps, ought not to have any doubts about whether or not the table is made of gorgonzola; however, *given that he has them* it would be irrational for Cohen to come to believe (5); in Cohen's first person case, his doubts can play an epistemically important role.

Therefore, dogmatist is not committed to the claim that one's total epistemic state (the totality of what one doubts, believes, knows, is warranted in believing and so on) cannot affect whether or not one warrantedly believes something on the basis of one's senses. The dogmatist claims that having a perceptual experience as of p gives one defeasible warrant to believe p. The *prima facie* character of the warrant plausibly transfers to any belief one subsequently bases solely on p; beliefs inferred from warranted perceptual beliefs will at best be warranted in a defeasible manner. It is compatible with a belief being *prima facie* warranted by one's perceptual experiences that other features of an agent's psychology could render that belief to be unwarranted, all things considered. In particular, if one doubts, worries about, or believes the negation of p, these doubts and/or contrary beliefs could undermine or rebut the perception-based warrant in such a way that all things considered one is not warranted in believing p.[22]

[22] One may wonder whether or not a simply doubt is enough to undermine ordinary perceptual warrant. If it cannot, then Cohen's objection does not get off the ground because even if he *does* doubt (for no good reason) that the table is red, or that it is normally lit, we should conclude he is still warranted in believing the table is red. In order to do justice to Cohen's argument I have simply assumed that the doubts of the kind he mentions are enough to undermine his perceptual warrant.

Cohen would remain unconvinced by the preceding considerations because he thinks that one can present the case in such a way as to not involve any doubts or questioning of the conclusion. He claims that he can "eliminate entirely the dialectical context" by having his son simply grant that Cohen knows that the table is not white with red lights shining on it, but he is just curious *how* Cohen knows this. If Cohen responds to this challenge by saying: "Oh, that's easy. It looks red, so it is red, so it is not white with red lights shining on it," Cohen maintains that this is problematic even though there is no obvious dialectical failure like question-begging going on (Cohen 2005, 420). This is because Cohen's son does not provide any reason to doubt the color of the table, nor does his son demand that Cohen provide an argument that convincingly establishes that the table is not white and illuminated redly. Furthermore, in this case neither Cohen nor his son has any doubts about Cohen's belief that the table is not white and illuminated redly. Cohen's son simply requests an explanation as to how Cohen knows what they both grant that he does know. Nonetheless, Cohen claims that there is still something problematic with the response he provides his son.

My points in the preceding paragraphs will not address this new case. Cohen does not have a defeating belief, nor any epistemically relevant doubts. However, the problem with Cohen's non-dialectical case is of a very similar kind. The claim that being immediately warranted in believing p on the basis of one's perceptual experience as of p does not require that one be able to produce non-question-begging arguments in favor of p derives from a more general commitment to avoid *over-intellectualizing* the requirements for perceptual warrant.[23] By reflecting on perceivers such as children,

[23] For a more complete discussion of this point see Chapter I Section 2.

higher animals, and unreflective adults, who all plausibly have warranted perceptual beliefs, it seems obvious that one need not know or be warranted in believing that the conditions on perceptual warrant are met in order to possess warranted perceptual beliefs. In particular, one need not know, warrantedly believe, or be able to articulate or defend (even to oneself) the correct epistemological explanation for one's possessing the warrant in question.

What Cohen's son is demanding in this amended case, is an epistemological theory that explains the nature of Cohen's warrant for his belief. Since Cohen's belief that the table is not white and illuminated redly is based on his perceptual belief that the table is red, an adequate answer to his son's question will involve specifying the nature of the warrant Cohen has for his perceptual beliefs. But any suitable theory of perceptual warrant ought to conclude that one need not understand, warrantedly believe, or know that very epistemological theory of perception in order to be warranted in one's perceptual beliefs or utilize them for simple inferences. Such a demand would yield the unpalatable result that *no one* (except for some fortunate epistemologists, perhaps) is warranted in their perceptual beliefs, because children, animals, most adults, and even many epistemologists (who believe false theories), do not know or warrantedly believe a fully worked out epistemological theory of perception. As such, we should not infer any conclusion about the conditions for warranted belief from Cohen's inability to answer his son's question about *how* he knows what he does. In reflecting on the case a conceptually sophisticated individual such as Cohen (or his son) could find himself confused and curious about the nature of the warrant that underwrites these beliefs. A similar confusion paired with curiosity presumably helped prompt some few

epistemologists to devote their professional lives to questions of this kind. I can sympathize with this attitude. However, we should not conclude on sober reflection that we actually *need* an answer to such questions in order to have warranted perceptual beliefs or warranted beliefs inferred from one's perceptual beliefs.

There is another consideration that strongly suggests that we should avoid drawing any conclusions about the conditions for warranted perceptual belief from Cohen's dialogs. If we do so, we will be committed to an implausibly strict constraint on warranted belief; so strict that any reasonable epistemic theory ought to reject it out of hand. As Markie (2005) points out, even if we stipulate that the father in Cohen's easy knowledge cases knows that his vision is reliable, this will not fully address his son's question. On a natural interpretation of his query, Cohen's son is not calling into question the *general reliability* of Cohen's vision, he is instead raising the possibility that *in this particular case* Cohen finds himself in a deceptive environment. One would think that Cohen's son would be equally unconvinced if his father were to protest as follows: "Look son, I know that my color-vision is reliable. The table appears to me to be red, and since my color-vision is reliable, it is red." "Reliable or not," his son should respond, "What I want to know is whether or not you are deceived *right now*. That you are ordinarily not deceived is of no help to me. I just want to avoid buying a white table!"

What this revision seems to show is that the problem, if there is one, is not special to dogmatism, reliabilism, or even any theory that denies WR. The demands of Cohen's son suggest a constraint with far wider scope than either Cohen or Markie acknowledge. In fact, it seems as if an analogous case can be generated for any theory that denies the following principle:

> **(UC):** In order for S to be warranted in believing p in a particular context C, S must be able to rule out any possibility that S is deceived or in error *in C.*

If there is even the possibility of deception in the particular case, then no matter what other epistemic virtues Cohen can trot out for his son to show that he is warranted in believing that the table is red, his son can always claim that he is uninterested in his father's *general* epistemic competence, what he wants to know is whether or not he has made an error or has been deceived *in this particular case*. If appeasement of Cohen's son is meant to have implications for our putative perceptual warrant in ordinary cases, it seems that UC captures the condition that we must meet in order to possess warrant.

But nothing about the case or Cohen's commentary on it does anything to lend support to UC, and UC amounts to a straightforward denial of one of the dogmatist's fundamental claims: namely that one's possessing perceptual warrant is compatible with error or deception in particular cases. There is a wide agreement among epistemologists that there is nothing objectionable about this feature of dogmatism.[24] It is an undeniable fact that, through no fault of our own, nor malfunction of our perceptual systems: sometimes our perceptual experiences do not accurately reflect the way the world is. Whether as a result of perceptual illusion, momentary malfunction, hallucination, or any number of other causes, sometimes we have non-veridical perceptual experiences. The possibility of global deceptions such as a brain-in-vat scenario suggests that we can never entirely rule out the possibility of deception or error in any particular case. The dogmatist acknowledges these points and asserts that, nonetheless, we can be warranted in believing

[24] Of course the agreement is not universal, McDowell (1996) is one such dissenter, as is Unger (1979). Cohen, however, gives no indication that he means to object to this element of the dogmatist position. Indeed, he seems to explicitly endorse it (Cohen 1988).

things on the basis of our perceptual experiences. UC, however, seems incompatible with this respectable feature of dogmatism. If UC is true we must be able to convincingly rule out the possibility of error in each particular case in order to possess a warranted perceptual belief belief. The appropriate response, therefore, is to reject UC rather than any theory it purports to undermine, and this seems entirely appropriate.

6. Wright's Simple Elevation Hypothesis

i. Two Levels of Explanation

Wright (2007) begins his most fully developed objection to dogmatism by discussing what he calls "justificational triads" which consist of (I) evidence for p, (II) the proposition p itself, and (III) an *authenticity condition*. (Wright 2007, 29) The third element of the triad is a proposition such that if one doubted its truth, then one could not rationally believe p on the basis of one's evidence for p. Wright gives several examples of such triads, but for our purposes here, discussion of one of these is sufficient:

> **(I)** S's experience as of a red cube.
> **(II)** This is a red cube.
> **(III)** My visual system is trustworthy and functioning properly. (Wright 2007, 40)

All parties in the debate can concede that if one were to doubt (III) of the above triad, one could no longer rationally believe (II) on the basis of (I).

Wright distinguishes between questions of whether a belief that p is warranted, and questions of whether a subject S is *in a position to claim warrant* for a belief that p.

> I want to understand the claimability of a warrant to be what is at issue when, for example, a philosopher feels that one has not been given everything one needs to address scepticism about the external world, say, merely by impressive arguments

> — if any such there be—that knowledge can be constituted by reliably generated true belief...the residual dissatisfaction with the externalist suggestion as a response to scepticism is that it merely points to a congenial possibility: nothing has been offered to put us in position to claim that it, rather than one of the many contrasting uncongenial sceptical scenarios, actually obtains (Wright 2007, 30-31).

Given the manner in which Wright marshals this distinction against dogmatism, I believe that we should take his language here quite literally. I take it that, for Wright, being in a position to claim warrant for p at least entails that the subject (if she is so inclined) be capable of rationally and correctly claiming warrant for p in a dialectic, particularly in response to challenges from an interlocutor.[25] Wright formulates two questions which one might ask of the justificational triads mentioned above:

> **Level-1:** What independent epistemic relation must S bear to (III) if having (I) is to give her warrant for (II)?
>
> **Level-2:** What independent epistemic relation must S bear to (III) if, when (I) does give her warrant for (II), she is *to be in position to claim* a warrant for (II)? (Wright 2007, 32)

Wright makes strong claims about the epistemic importance of the level-2 question.[26] However, all he needs for his main argument to have force against dogmatism is the

[25] I argue for this claim in the next sub-section. There I contend that unless we understand the condition to have this entailment the case Wright considers will not generate any conclusions concerning the dogmatist's position. Wright does not characterize "position to claim" in any more rigorous a fashion than what is given in the passage above, nor does he indicate what sorts of conditions must be met in order to achieve the state of being in a position to claim warrant for p. However, the primary problem with Wright's argument is not the ambiguity of his terminology, so I suggest that we leave the characterization as it stands, and proceed to the manner in which Wright deploys it.

[26] Wright argues that the level-2 question is a "cardinal objective of the traditional epistemological project, properly conceived" (Wright 2007, 31). He claims that any account of warrant which does not provide an acceptable response to the level-2 question

relatively weak claim that the level-2 question is an important epistemological question, and that it would be a mark against dogmatism if the theory entailed an unacceptable response to this question. I grant this claim for the purposes of this discussion.

Before proceeding to Wright's argument that the dogmatist is committed to a poor response to the level-2 question, it is worth pausing to reflect on the difficulties besetting such a project. Claims about whether a belief is warranted are distinct from claims about higher-level epistemic relations to that warrant. Being in a position to claim warrant is surely one of the latter. I do not deny that some account of these higher-level epistemic attitudes must be provided by a complete epistemology. Dogmatism, however, is characteristically a theory that concerns the warrant furnished by perceptual experience for one's basic perceptual beliefs. In general, claims about lower-level epistemic positions do not entail claims about their higher-level correlates. For example, claims about knowledge do not generally entail claims about knowledge of knowledge, claims about warrant to believe p do not entail claims about warrant to believe that one is warranted in believing p, and so on. It would be surprising, therefore, if the dogmatist were implicitly committed to a view about something that he explicitly eschews taking a position on.[27]

(he targets *externalism*) has failed to address the deep issues that concern us regarding skepticism and perceptual warrant generally (Wright 2007, 36). I believe that such claims are mistaken, but addressing them is beyond the scope of the present discussion. Wright's argument may proceed with the weaker claim I grant above.

[27] Again, for admirable discussion of this and related issues see Alston (2008)

ii. The Simple Elevation Hypothesis

Wright, perhaps anticipating the concerns discussed at the end of the preceding sub-section, provides a principle that serves to link the level-1 and level-2 questions in the appropriate way. He calls his proposed principle the Simple Elevation Hypothesis (SEH):

> **SEH:** In general, whatever conditions C (a priori necessarily) confer warrant upon S for acceptance of p, it will suffice to put S in position to claim warrant for p that she be in position to claim that C are met (Wright 2007, 33).

In other words, if a certain set of conditions C confer warrant for a subject S to believe a proposition p, and if S is in a position to claim that C are met, then S will be in a position to claim warrant for a belief in p. I will evaluate the plausibility of this principle in due course; for the moment, let us see where its acceptance by the dogmatist takes us with regards to the level-1 and level-2 questions. The dogmatist's answer to the level-1 question above seems clear: the only independent relation one must bear to the claim that one's perceptual system is functioning normally in order for the experience as of red cube to warrant one in believing that there is a red cube (on the basis of one's visual experience as of a red cube) is that one has no reason to doubt that one's visual system is functioning normally. If the dogmatist accepts SEH, then he will be committed to an answer to the level-2 question as well. SEH claims that one is in a position to claim warrant for a proposition if one is in a position to claim that sufficient conditions on a warranted belief have been met. It follows that the dogmatist is committed to the following answer to the level-2 question: in order to be in a position to claim warrant for a basic perceptual belief p, one need only have the experience as of p, no reason to doubt that one's visual system

is functioning properly and be in a position to claim that both of those conditions are met (Wright 2007, 34-35).

Wright regards this answer to the level-2 question as unacceptable. He argues by use of imagined dialogs in which a subject claims warrant for a proposition in the face of questions from an interlocutor. I reproduce a version of the most pertinent dialog:

> [Interlocutor] 'Why do you believe that there is a red cube here?'
>
> [Subject] 'Because it looks to me to be the case that there is a red cube here and I have no reason to regard my visual system as defective.
>
> [Interlocutor] 'So: you take it that it is reasonable to believe that your visual system is likely to be trustworthy unless there is evidence that it may not be functioning properly?'
>
> '[Subject] *I do not need to have a view about that*. I repeat: it looks to me that there is a red cube and I have no reason to doubt that my visual system is working properly. I therefore claim warrant that there is a red cube' (Wright 2007, 39).

Let us stipulate that in the above dialog we are dealing with a conceptually sophisticated subject, and that the subject in question is in a position to claim that the dogmatist's conditions on warrant for basic perceptual beliefs are met. In other words, the subject is in a position to claim both that he is having a perceptual experience of a red cube, and that he has no reason to doubt that his visual system is working properly. If we accept SEH, it seems that the dogmatist is committed accepting the subject's final response of the dialectic as acceptable. The subject is in a position to claim that the correct conditions (according to the dogmatist) are met. It follows from SEH that he is in a position to claim warrant, and so there can be nothing wrong with his claiming warrant in response to the interlocutor. Wright contends that this is an undesirable result:

> ...there is a very strong sense that the last response is irrational, or epistemically irresponsible. After his first answer, and once the issue is raised, the speaker does need to take the stance that it is reasonable to believe that his visual system is

> working properly unless there is evidence to the contrary — or back off his claim to visual warrant for p" (Wright 2007, 40).

In short, taking no stance on the trustworthiness or functioning of one's visual system is irrational if one is claiming warrant to believe that there is a red cube here on the basis of one's visual experiences. In order to claim warrant in the face of the interlocutor's questions one *must*, Wright thinks, rationally believe that one's visual system is trustworthy and functioning properly (Ibid.). We can grant provisionally that there is something amiss with the final response in the above dialectic, and since, as I have understood it, being in a position to claim p is essentially tied to what one can properly and rationally do in response to challenges from an interlocutor, we can grant that something is wrong with the subject's claim to warrant. Wright concludes that the dogmatist is committed to an unacceptable response to the level-2 question (since they must accept SEH). If this were so, this would be a mark against dogmatism.

It is now clear why we must understand the condition on a subject's being in a position to claim warrant in the manner in which I have urged. Wright can only argue as he does if a subject's being in a position to claim warrant for p entails that the subject can (if so inclined) rationally and correctly claim warrant in a dialog with an interlocutor. If we do not understand the condition in this way the dogmatist could simply accept SEH, accept that the subject in the above dialog is in a position to claim warrant, while *also* accepting that there is something inappropriate about the subject's final response. The subject would be in a position to claim warrant, but would be unable to do so properly and rationally in the imagined dialog. The dialog would, therefore, indicate nothing about the plausibility of the dogmatist's proposed answer to the level-2 question. To be

sure, the breakdown between being in a position to claim warrant for p and being able to rationally do so in the dialectic would need to be accounted for, but since nothing precludes such a mismatch, there is no reason to think that an explanation could not be produced. One could, for instance, argue that the subject's failure in such a case is a pragmatic or dialectical failure: e.g. while it is true that the subject is in a position to claim warrant, doing so is dialectically inappropriate because the claim to warrant is precisely what is at issue in the discussion. There is more to be said along these lines, but there is no need to pursue this matter further in this context; it is clear that Wright is not interested in arguing along these lines. He immediately infers that the inappropriate nature of the subject's final response indicates that he is *not* in a position to claim warrant, and hence that the dogmatist's answer to the level-2 question entailed by SEH is incorrect (Wright 2007, 40). I see no charitable recourse but to conclude that, as Wright understands it, a subject's being in a position to claim warrant for p entails that the subject can rationally and correctly claim said warrant in response to an interlocutor's challenges. If we understand the condition such that it does not have this entailment, the dogmatist has an obvious line of response that Wright does not address.

iii. Against the Simple Elevation Hypothesis

The most straightforward way to undermine Wright's criticism is to challenge the principle which he utilizes to force the dogmatist's hand; namely SEH. Wright himself proposes what I take to be a good reason to reject this principle:

> Someone might demur from [SEH] on account of the following gap: conditions C may obtain, so that S has a warrant for p, and S may be in position to claim that C

> are met, and yet in no position to claim warrant for p because she doesn‘t know that C *are* warrant conferring for p (Wright 2007, 33).

This seems correct to me. SEH does not include any requirement that the subject know, be warranted in believing, or be in a position to claim that C are warrant conferring conditions. Without such a requirement, the principle seems straightforwardly false. Let us consider a subject, Haley, who is cognitively limited in such a way that she does not have the capacity to consider propositions about warrant. According to the dogmatist's account of perceptual warrant, Haley may still meet the conditions sufficient to be warranted in believing p. Suppose that she meets these conditions, that she has just enough cognitive complexity to consider and believe propositions concerning the relevant conditions, and that Haley is in a position to claim that the conditions are met.[28] It seems clear in this case that since Haley cannot even entertain propositions concerning warrant, she cannot be in any position to claim that she is warranted. Thus, Haley is in a position to claim that the conditions for a warranted belief are met, but she is *not* in a position to claim that she is warranted.

Wright concedes that cases like this bear against SEH (Wright 2007, 34). However, he does not think that the dogmatist can appeal to cases like Haley's in order to argue against SEH. To show this he considers a different case:

> Sophie is a six-year-old child who forms beliefs about her immediate environment on the basis of her perceptions. She believes that her room is untidy, that her cat

[28] The possibility of a case like this will depend on the conditions for a warranted belief. In the case of dogmatic warrant for perceptual belief, it seems that such a case is possible. A small child could, seemingly, be in a position to claim that things look to her to be a certain way and that there is no funny business going on (like someone trying to trick her) without also being capable of understanding epistemic concepts such as *warrant*.

> is sleeping in its bed, etc. on the basis of her visual experiences of these things. Sophie also has a rudimentary understanding of the word "know" and when asked she claims that she *knows* where her cat is and that she *knows* that her room is untidy. (Wright 2007, 34)[29]

Wright draws three conclusions about this case: first, he grants along with the dogmatist that it seems correct to say that Sophie knows that her room is untidy and the cat is in the bed. He also claims that it seems correct to say that Sophie also knows the claims she makes about her own knowledge. He concludes that:

> These [knowledge claims] are effectively claims of warrant but they seem just as justified and rational, in the circumstances, as the claims they embed. It would be quite out of keeping with the general tenor of dogmatism to insist that they are not so — that warrant for them requires in addition a sophisticated piece of epistemological knowledge (Ibid.).

It should first be noted that dogmatism says nothing whatsoever about the conditions that a subject must meet in order to make warranted claims about knowledge, warrant, or any other epistemic state. Perhaps some dogmatists might be willing to accept that Sophie is warranted in her higher-level assertions, but it is entirely consistent with the view as stated to deny this claim and to impose more strict conditions on warrant for beliefs about one's epistemic status. It is, therefore, open to the dogmatist to simply deny Wright's claims in the above passage.

However, let us grant that Sophie does have the higher-order knowledge in question in order to ascertain whether or not Wright's argument presents a problem for dogmatists inclined to accept his stipulations. Does her case show that such a dogmatist cannot appeal to cases like Haley's in order to deny SEH? In other words, does Sophie's

[29] I ignore the fact that these beliefs are probably not perceptually basic. The same points could be made with my example of a red cube.

case show that being in a position to claim that the conditions on dogmatic warrant are met are sufficient for her to be in a position to claim warrant or knowledge? I think not. The dogmatist can claim that whatever mastery that Sophie's competence with the concept *knowledge* is playing a crucial role in the appropriateness of her claim, and hence that simply knowing that things appeared to her a certain way, and believing that she had no good reason to believe she was being tricked, might have failed to put her in a position to claim knowledge. Reflection on Haley's case above suggests that if Sophie were to lack mastery of the concept of knowledge and were still able to report on how things seem to her and whether or not she is being fooled in some way, might still fail to put her in a position to claim warrant. The dogmatist may argue that Wright's "additional piece of sophisticated epistemological knowledge" is precisely what Sophie possesses and Haley lacks: namely an understanding and ability to deploy the concept of *knowledge.* Sophie's case fails to show that the dogmatist cannot make use of cases like Haley's to argue against SEH. The dogmatist may reject SEH for precisely the reasons that Wright himself suggests.

My discussion, however, suggests a further avenue of attack available to Wright. In a rather speculative discussion Wright suggests that even if SEH is not true of a subject such as Haley who is "not enough of an epistemologist" to possess the relevant concepts or understand the epistemic relationships between her various beliefs, it may still be true of subjects who *do* possess this understanding (Wright 2007, 35). Thus SEH may be false for epistemic innocents such as Haley, but still be true for subjects who are capable of reflecting on their epistemic status.

> If this were the situation, the Simple Elevation Hypothesis would remain unchallenged. The mismatch between the dogmatist's answer to [the level-1

> question] and the answer it enforces, in accordance with the Simple Elevation Hypothesis, to [the level-2 question], would be explained by the consideration that the very conceptual sophistication required by the practice of self-ascription of warrant of the kind illustrated by the dialogue somehow 'ups the ante': that it raises the standards for having warrant in the first place. So the liberal answer to [the level-1 question] would be correct only for relatively unsophisticated subjects; and no answer, liberal or conservative, would be generally correct. (Wright 2007, 35)

In other words, Wright might be willing to concede that dogmatism gives the correct story for epistemically unsophisticated subjects such as Haley, while still maintaining that it fails to provide the correct account of perceptual warrant for subjects with the proper degree of conceptual sophistication. SEH may be false of epistemic innocents such as Haley, yet may still be true of subjects capable of entertaining propositions concerning their own epistemic states. Wright's argument would, therefore, fail generally, but still saddle the dogmatist with an unpalatable result for epistemically sophisticated agents.

SEH fails for more sophisticated agents as well and Wright himself may constitute a counter-example to the restricted principle. Let us say that Wright (with his current philosophical views about perceptual warrant) perceives a red cube in ordinary circumstances. He is in a position to claim that it appears to him that there is a red cube before him and that he has no reason to doubt that his perceptual system is representing the scene before him veridically. Is Wright *thereby* in a position to claim warrant for this belief? It is clear that he is not. Wright believes that the satisfaction of these conditions is insufficient for a warranted belief based on his visual experience. There would be something irrational or epistemically irresponsible in Wright's claiming warrant for a belief for which he has not met all of what he takes to be the necessary warrant conferring conditions. The dogmatist can claim in such a case that Wright is in a position

to claim that he meets the conditions necessary for possessing dogmatic warrant to believe that there is a red cube in front of him, but because of his beliefs about warrant is not, thereby, in a position to claim that he is so warranted.[30] Conceptually sophisticated individuals with contrary beliefs about the conditions under which beliefs are warranted will be counterexamples to SEH.

The most obvious way for Wright to respond to the above arguments would be to revise SEH to include the condition that the subject also be in a position to claim that the relevant conditions are *warrant conferring* in order to be in a position to claim warrant. Thus he could propose the following:

> **SEH':** In general, whatever conditions C (a priori necessarily) confer warrant upon a thinker S for acceptance of p, it will suffice to put S in position to claim warrant for p that she be in position to claim that C are met and that she be in a position to claim that C are warrant-conferring.[31]

SEH' seems to have some antecedent plausibility and is not susceptible to the same kinds of counter-examples I have marshaled against SEH. Furthermore, SEH' does impose some commitments on the dogmatist with regards to Wright's level-2 question. However, it is not clear that these commitments are problematic. Consider a dialog similar to the one Wright constructs, but in which the subject meets the antecedent conditions of SEH':

[30] I leave it open what the dogmatist should say about whether or not Wright is warranted in this case. It may be that possessing false philosophical beliefs about the conditions for perceptual warrant could undermine the warrant for one's basic perceptual beliefs. However, it is consistent with the general spirit of dogmatism that such philosophical errors do not undermine one's basic perceptual warrants (or at least usually don't). In forming typical perceptual beliefs, Wright's views about the nature of perceptual warrant are plausibly not epistemically relevant to the formation or sustaining of his perceptual beliefs. I remain neutral on this issue here.

[31] One could also revise the principle to include that the subject know or warrantedly believe that the C are warrant conferring. My remarks concerning SEH' apply to these variations of the principle as well.

> [Interlocutor] 'Why do you believe that there is a red cube here?'
>
> [Subject] 'Because it looks to me to be the case that there is a red cube here and I have no reason to regard my visual system as defective.
>
> [Interlocutor] 'So: you take it that it is reasonable to believe that your visual system is likely to be trustworthy unless there is evidence that it may not be functioning properly?'
>
> [Subject] '*I don't need to address that question right now.* If I have no reason to doubt the proper functioning of my visual system, then I am warranted in believing what it represents to me. These are the conditions I must meet in order to be warranted in this sort of belief (and I am in a position to claim this). I repeat: it looks to me that there is a red cube and I have no reason to doubt that my visual system is working properly. I therefore claim warrant for beleiving that there is a red cube.'

There does not seem to be anything obviously problematic with the interlocutor's responses in this dialog. If *anything* will be sufficient to put a subject in a position to claim warrant, being in a position to claim that a certain set of conditions are warrant conferring and being in a position to claim that those conditions are met surely should suffice. If a subject is in a position to claim that a particular epistemological theory of warrant is the correct one, and that she has met all the conditions said theory imposes on warrant, it is unclear what else would be required to put her in a position to claim that she is so warranted. I should note that this concession does not commit the dogmatist to any view about the conditions a subject must meet in order to put her in a position to claim that a set of conditions are warrant conferring or that they obtain. Given SEH', all the dogmatist is committed to is that *if* a subject finds themselves in a position to make these claims, then she will also be in a position to claim warrant. We have been given no reason to suppose that there is anything objectionable about this commitment.

It may be that in order to be in a position to claim that the conditions in question are warrant conferring that one would need to address the interlocutor's final question.

That is, in order to be in a position to make the dogmatist's claim about the conditions under which subjects may be warranted in their perceptual beliefs, one may be required to offer an argument establishing that trusting one's visual system in the absence of defeating or undermining evidence is reasonable. This amounts to a demand that the dogmatist present an epistemological explanation of precisely why the conditions he proposes are warrant conferring. Burge (2003) and others have attempted to develop such explanations. I discuss the virtues and vices of these various approaches in the third and fourth chapters of this dissertation. However, an objection to dogmatism along these lines must either present considerations which suggest that no such explanation can be provided, or argue that all the views on offer fail in some way or other. Wright presents no such argument. Given that SEH is false and that the proposed revision of the principle, SEH', does not entail clearly counterintuitive results, I conclude that these principles do not present a substantial problem for dogmatism.

7. The Bootstrapping Objection

i. The Bootstrapping Procedure

Roxanne wonders whether or not her color-vision is reliable and setting some time aside one day conducts an investigation on the subject. In order that her investigation be as rigorous as possible she discounts any evidence or reasons she may have previously had concerning the reliability of her color-vision and resolves to start from scratch. Roxanne has arranged that a series of randomized color cards be displayed one after the other. As each card is displayed she writes the following premises down in her experiment log and forms the corresponding belief:

(1) My visual experience is as of a red card.
(2) The card is red.
(3) My visual experience is as of a red card, and the card is red.
(4) Therefore, my visual color experience matched the actual color of the card on this occasion

Roxanne arrives at (1) by introspecting on the character of her visual experience. She believes (2) solely on the basis of her visual experience as of a red card. Roxanne arrives at (3) by competently deducing it from (1) and (2), and subsequently carries out the deductive inference to (4). Roxanne carries out this procedure for each of a long series of cards of varying colors. At the end of this series, she competently deduces the following from the results:

> **No Errors:** My color experiences have matched the actual color of each of the many cards that I have viewed.

Given No Errors and a sufficiently large number of trials, Roxanne concludes that her color-vision is reliable. Call the pattern of reasoning that Roxanne has carried out *bootstrapping*.[32] It is almost universally accepted that Roxanne cannot come to be warranted in believing that her color-vision is reliable by bootstrapping her way to this conclusion.[33] Roxanne has made no independent checks to confirm that her appearances are accurate in any of the individual trials of her experiment. She has simply taken her visual experiences at face value and concluded that her color-vision is reliable on the basis of these "trials." Surely one cannot come to know that a source of information is reliable in such a manner. An epistemic view *sanctions bootstrapping* if the view entails

[32] This presentation of the bootstrapping procedure is a mash-up of the presentations of White (2005) and Cohen (2002).

[33] Van Cleve (2003) is an exception.

that Roxanne's procedure is sufficient to warrant her belief that her color-vision is reliable. Because bootstrapping intuitively does not warrant Roxanne in believing that her color-vision is reliable, any view that sanctions bootstrapping is objectionable for this reason.

The dogmatist is committed to claiming that Roxanne is warranted in believing (2), so long as she does not have defeating evidence for her belief in the color of each card. However, this does not suffice to establish that the dogmatist must sanction bootstrapping. The dogmatist makes no claims about the conditions necessary for warrant generally. Dogmatism only weighs in on the conditions under which perceptual beliefs are warranted. The view does not entail anything about the conditions under which beliefs may be warrantedly inferred. Strictly speaking, dogmatism is consistent with a radical Humean skepticism about induction that maintains that no such inference is ever warrant-preserving.[34] Indeed, dogmatism is compatible with claiming that no inference, deductive, inductive or otherwise, ever preserves warrant. Thus, one cannot simply appeal to entailments of the dogmatist's claims in order to show that the view sanctions bootstrapping.

What the objector ought to argue is that if dogmatism is supplemented with plausible epistemic principles the conjunction of dogmatism and these principles will entail that Roxanne, contrary to intuition, is warranted in believing that her color-vision is reliable. This strategy should invoke principles whose denial would be counted as a severe theoretic cost; the more plausible the principles that are used to commit the dogmatist to sanctioning bootstrapping, the stronger the argument against dogmatism will

[34] This example is due to a very helpful conversation with Michael Rescorla.

be. There are three kinds of non-perceptual belief-forming methods that Roxanne utilized in her procedure. First, she used introspection in forming her beliefs about the character of her visual experiences. Second, she employed a series of valid deductive inferences in order to eventually infer No Errors. Finally, she employed a non-deductive inference in order to conclude, from No Errors, that her color-vision was reliable. If each of these inferences can be supported by plausible epistemic principles, then it would seem that the dogmatist would be faced with the unpalatable dilemma of either sanctioning bootstrapping or rejecting one or more of these principles.

The dogmatist is committed to claiming that Roxanne's beliefs in the color of each of the cards are warranted since they are based on undefeated perceptual experiences.[35] In what follows, I simply take for granted that Roxanne's introspective beliefs are warranted. The objector requires principles that entail that if, in each trial, Roxanne is warranted in believing (1) and (2), she will also be warranted in believing each of the subsequent premises in each trial, No Errors, and her conclusion that her color-vision is reliable. In each trial Roxanne competently deduces (3) and (4) from her introspective and perceptual beliefs. She deduces No Errors from the conjunction of (4)'s from each trial. Thus, if the dogmatist denies that Roxanne is warranted in any of her inferences leading up to No Errors, she will be committed to denying that warrant is closed under competent deduction:

> **Closure of Warrant (CW):** If S is warranted in believing p, p entails q, and S competently deduces q from p, then S is warranted in believing q.

[35] If one wishes to endorse dogmatism of an externalist flavor, the case must be described in such a way that Roxanne meets the externalist constraints on perceptual warrant, but is unaware (or does not take into account in her reasoning) that she meets these constraints. The argument proceeds as otherwise stated.

Closure principles such as CW are controversial. Though it can seem obvious that something along these lines must be true, it is quite difficult to formulate an exceptionless statement of the plausible idea that competent deduction is always warrant-preserving.[36] Still, the core idea behind CW is plausible enough that being forced to deny the principle is rightly reckoned a severe theoretic cost. Furthermore, even if CW is false, it is certain that competent deduction *usually* preserves warrant, and it is quite difficult to see what could be going wrong in any of Roxanne's deductive inferences. The entailment in each case seems obvious, and there don't seem to be any obvious features that would lend themselves to a plausible explanation of why they fail to preserve warrant. As such, defending dogmatism by rejecting CW is a problematic strategy, and one I will not pursue here.

If the dogmatist accepts CW, she is committed to claiming that all of Roxanne's beliefs up to and including No Errors are warranted. However, the inference from No Errors to the claim that her color-vision is reliable is not deductive. That one's color-vision produced a large series of veridical representations does not entail that one's color-vision is reliable. It is possible that the series of veridical representations is not the result of reliability or competence on the part of one's perceptual system, but an extraordinary run of luck. One's perceptual system may provide veridical color representations only half of the time, but by pure chance have a long run of correct answers. Thus, CW alone does not commit the dogmatist to accepting that Roxanne is warranted in believing her color-vision is reliable.

[36] For a useful survey as well as a case that no such principle is exceptionless see Hales (1995).

Roxanne's final non-deductive inference is typically supported in one of two ways. Cohen (2006) argues that Roxanne could use *enumerative induction* to support the claim that her color-vision is reliable. Given a large enough series of trials, Roxanne's inductive base could be considered large enough to ground a well-supported inductive inference to the claim that her color-vision will continue to produce substantially more veridical color representations than non-veridical ones, and hence that her color-vision is reliable. On the other hand, White (2006) claims that Roxanne's final inference is best seen as *abductive*. Supposing that Roxanne has carried out this procedure through many trials, the best explanation of the success of her color-vision in such a large number of cases is that her vision is reliable in its representations of colors. That she simply got extraordinarily lucky in getting things right so many times seems a far poorer explanation of the evidence available to her, hence she is warranted in concluding that her color-vision is reliable.

If there is nothing epistemically suspect with Roxanne using either enumerative induction or an abductive inference to draw her final conclusion, then the dogmatist is committed to the unpalatable consequence that Roxanne can bootstrap her way to a warranted belief in the reliability of her color-vision. In what follows I argue that Roxanne's non-deductive inference, whether understood inductively or abductively, is not plausibly warrant-preserving, and that the diagnosis of why this is the case is consistent with dogmatism.

ii. Bootstrapping and Self Confirmation

Let us to attempt to articulate precisely what is so problematic about Roxanne's bootstrapping inference.[37] This diagnosis will be instrumental in showing how the dogmatist can avoid sanctioning the procedure. Consider a case structurally similar to Roxanne's in which the reasoning does not seem to be objectionable. Enid is also curious about her color-vision, and undertakes an investigation similar in many ways to Roxanne's. Enid reasons as follows:

(1) My visual experience is as of a red card.
(2) The card is red.
(3) My visual experience is as of a red card and the card is red.
(4) Therefore, my color-vision was accurate on this occasion.

After repeating (1)-(4) for a large series of different cards of different colors, Enid infers:

> **No Errors:** My color experiences have matched the actual color of each of the many cards that I have viewed.

Given No Errors, and a sufficiently large number of trials, Enid concludes that her color-vision is reliable. Enid, unlike Roxanne, does not base her belief in each of her (2)'s on

[37] Van Cleve (2003) has suggested that there is nothing wrong with Roxanne's procedure. He argues that an epistemology that does not sanction bootstrapping is doomed to endorse skepticism and attempts to show how purported attempts to condemn bootstrapping while avoiding skepticism fail to do so. While I am inclined to agree that endorsing bootstrapping is a better alternative than skepticism, I am more optimistic about our chances of finding some way to accommodate both intuitions than is Van Cleve. In particular, I take the suggestion that bootstrapping could be an epistemically allowable method of testing the reliability of one's faculties to be ridiculous enough, that sanctioning bootstrapping ought to be a last resort. I am neutral as to just how bad this result would ultimately be. If the dogmatist could not avoid sanctioning bootstrapping, what would be required is a cost/benefit analysis along the lines of Van Cleve's discussion. One would have to weight the benefits of dogmatism in dealing with other problems against the apparent absurdity of sanctioning bootstrapping. My argument can be seen as maintaining that no such extreme measures need be taken.

her color-vision. Instead, she refers to an independently verified answer key that tells her what the color of each respective card is. Let this independent verification take whatever form the reader deems necessary in order to make it obvious that the key itself is reliable and Enid is warranted in believing it to be such. I take it that few would deny that Enid comes to be warranted in believing that her color-vision is reliable in this way.

Enid and Roxanne's procedures share structural similarities. In discussing the differences between them, it will be useful to stipulate some terminology. Enid and Roxanne both make use of what I will call *Track-Record arguments*.[38] A Track-Record argument attempts to establish the reliability of a belief-forming method by comparing the ratio of true to false beliefs that the method generates. Call each case in the sample set a *trial*. A trial consists of noting what the method represents as being the case, and a statement of the way things actually are (i.e. premises (1)-(4) of Enid and Roxanne's procedures). If it can be shown that a belief-forming method produces many more true beliefs than false ones, this can provide strong support for the conclusion that the belief-forming method is reliable. Call the statements of the way things actually are, (i.e. the premise (2)'s in Enid's and Roxanne's procedures), the *confirmation clauses* of the Track-Record argument.

The only difference between Enid's and Roxanne's investigations is the source of the warrant for their respective confirmation clauses. Checking an independent and reliable source supports Enid's confirmation clauses; Roxanne, on the other hand, bases her belief in her confirmation clauses on her visual-experiences, the very faculty currently under investigation. The most straightforward diagnosis of Roxanne's failure is

[38] This term is due to Fumerton (1995) who presents an argument against reliabilism that is very similar to that presented by Vogel (2000).

that she has not done what Enid has done, namely utilized an independent check on the information supplied by her visual system. Intuitively, in order to know that one's color-vision is reliable a subject must do more than just look at the cards; she must use some other means to confirm that the color of the card matches the color the her experiences represents it to be in each trial. Enid satisfies this condition while Roxanne does not, thus Enid is warranted in believing her color-vision is reliable and Roxanne is not.

Let us present the diagnosis of the malignancy of bootstrapping as a principle:

> **No Self-Confirmation (NSC):** One can come to be warranted in believing that a belief-forming method M is reliable by use of a Track-Record argument only if one does not base one's belief in one's confirmation clauses solely on M.[39]

According to NSC, Roxanne cannot rely solely on her color-vision in order to support the confirmation clauses of her Track-Record argument. But, in basing her belief of the confirmation clauses on her visual experiences of the cards, this is precisely what Roxanne does. NSC implies, therefore, that Roxanne will not be warranted in believing that her color-vision is reliable.

Since dogmatism is neutral with regards to the conditions necessary for a strong Track-Record argument, it is entirely consistent for the dogmatist to endorse NSC as a diagnosis of what is wrong with Roxanne's procedure. In order to avoid sanctioning bootstrapping, then, the dogmatist may simply invoke NSC in order to block Roxanne's final inference from No Errors to the conclusion that her color-vision is reliable. The dogmatist (assuming CW) is committed to the claim that in each trial Roxanne is

[39] This principle could be an instance of a more general principle that applies to many different argument-types in addition to Track-Record arguments. However, this presents no complications for my argument unless it can be shown that this more general principle conflicts with dogmatism. I will argue in what follows that NSC does not.

warranted in believing (1)-(4), as well as No Errors. However, the dogmatist has no explicit commitments to the conditions under which non-deductive inferences such as Roxanne's are warrant preserving. Thus, even if the dogmatist is committed to sanctioning all of Roxanne's deductive inferences, the dogmatist can endorse NSC in order avoid sanctioning Roxanne's final non-deductive inference. If I am correct about my diagnosis of what is wrong with Roxanne's reasoning, then it seems that all that is necessary for the dogmatist to avoid sanctioning bootstrapping is for the view to be compatible with NSC.

Let us apply NSC to the two most popular means of supporting Roxanne's final inference.[40] Cohen (2002) and Vogel (2000) present Roxanne as utilizing *enumerative induction* in inferring her conclusion from No Errors, while White (2006) contends that the inference is best understood as *abductive*. Let us first examine Roxanne's procedure understood as an instance of enumerative induction. The case is stipulated to include the proper number of trials necessary in order for Roxanne to have a sufficiently large inductive base on which to ground her inference. However, the *size* of the inductive base is not all that is relevant. In particular, in constructing one's inductive base, one must avoid biases of various kinds. As NSC indicates, the bias in Roxanne's case has to do with the way in which she has confirmed, in each individual case, that the color her

[40] The epistemology of abductive and inductive inferences is a complex and controversial topic, one worthy of a treatise in its own right. My purpose here is not to provide a fully worked out framework of non-deductive inferences, but rather to propose what I take to be plausible explanations of the failure of Roxanne's procedure that are consistent with dogmatism. As I will make clear in the next section, little rides on whether or not the particular diagnoses I explore here are correct. So long as there is some plausible explanation for the failure of Roxanne's final inference, the dogmatist can invoke NSC to avoid sanctioning bootstrapping.

experience represented the card to be matched the actual color of the card. Roxanne's procedure relies on a correlation between her introspective judgments about her experiences, and the contents of those experiences. Notably, it does not rely on any correlation between her color-experiences and the actual color of the cards. Even if her color-vision were systematically unreliable, the correlation between the contents of her introspective judgments and those of her experiences would still hold, and Roxanne would still carry out her procedure just as described in the case.[41] Roxanne's procedure can tell her nothing about whether or not her color-vision is reliable because the results of her procedure do not covary with its reliability.[42] This diagnosis is entirely consistent with dogmatism. None of these claims entail that Roxanne's perceptual beliefs are unwarranted, or merely mediately warranted. Thus, the dogmatist is free to endorse this explanation of the failure of Roxanne's inductive inference in order to avoid sanctioning bootstrapping.

White's (2006) version of the argument has Roxanne employing an inference to the best explanation in order to infer her conclusion from No Errors.

---

[41] Thanks to Kevin Falvey for a helpful discussion and suggestions on these points.

[42] White considers and rejects a response on behalf of the dogmatist that maintains that the problem with Roxanne's procedure is that it could never *disconfirm* the hypothesis that her color-vision is reliable. If the impossibility of disconfirmability is the problem with Roxanne's final inference, the dogmatist is in serious trouble. As White points out, Roxanne's procedure also precludes the possibility of disconfirming No Errors (White 2006, 546-547). This diagnosis of the malignancy of bootstrapping entails, therefore that Roxanne is not warranted in believing even No Errors. But, as I have acknowledged, on pain of denying CW, the dogmatist is committed to claiming that Roxanne's belief in No Errors is warranted. My NSC-based diagnosis of Roxanne's failure avoids this unpalatable dilemma by explicitly ruling out her final inference. NSC, unlike White's disconfirmability principle does not entail that Roxanne's belief in No Errors is unwarranted.

> For what else could account for my long run of successful color judgments if not the reliability of my color-vision? It would be quite incredible for unreliable faculties to accidentally give the right results every time! It would be like someone guessing the outcome of a roll of dice many, many times in a row. We would have to suspect that there was something more going on to account for his success. As there is no other plausible explanation available to me in this case for my success at judging colors, I must gain some support for the reliability of my color-vision (White 2006, 546-547).[43]

White's contention is that the dogmatist must regard No Errors as a suitable evidence base for Roxanne's inference because the reliability of her color-vision is the best explanation for the long run of successful trials. But the dogmatist has no reason to accept this diagnosis. For it to be the case that a hypothesis best explains some evidence requires that one take into account *all of the relevant evidence*. NSC suggests that when one is attempting to determine whether or not some source of information is reliable by use of a Track Record argument, the source of one's support for the confirmation clauses is relevant. Roxanne's total evidence does not only include each of the premises she has written in her experiment log, it also includes the manner in which she came to believe each of those premises. In particular, NSC claims that the source of her confirmation clauses is something that she must take into account when she is deciding what best explains the truth of No Errors. But given that the support for her confirmation clauses is simply the testimony of her color-vision, it seems that the best explanation of No Errors is not the reliability of her color-vision.

To see why, it is useful to consider Enid's analogous abductive inference. This inference is plausibly strong because she utilizes an answer key that she has good reason to believe is reliable. Given that her color-vision has matched the key in each case (and

[43] Cohen (2006) endorses this diagnosis.

given a sufficiently large number of trials) the likelihood of this occurring merely by coincidence is very small. It is far more likely that the consistent matching between her color-vision and the independently verified and reliable answer key is due to the reliability of her color-vision. Applying White's diagnosis to Enid's case is apt: it would be unlikely in the extreme for Enid's color-vision to match the answer key in each case by sheer luck. The best explanation for this match, then, is that her color-vision is reliable.

Roxanne cannot rationally reason in this way. In each case she has merely matched her introspective judgment about her color-experience to the contents of the experience itself. That this was her procedure is relevant to an evaluation of what best explains the results of her trials. Roxanne can satisfactorily explain her results by the much more economical hypothesis that her introspective judgments about her perceptual experiences track the contents of those experiences.[44] This result is consistent both with the claim that her color-vision is reliable, and with its negation. Without further evidence or premises, there is simply no good reason to suppose that her introspective competence is best explained by any perceptual competence. As NSC states, Roxanne cannot simply ignore the source of her warrant for her confirmation clauses. In constructing a Track-Record argument as she does, the source of the warrant for her confirmation clauses constrains the abductive inferences that she can make from her premises.

Once again, this diagnosis of the failure of Roxanne's final inference is entirely consistent with dogmatism. It does not involve denying that Roxanne has immediate and defeasible warrant to believe the contents of her perceptual experiences. Rather, the

[44] Thanks to Kevin Falvey for a very helpful discussion on these issues.

problem rests in how Roxanne utilizes her warranted beliefs in her subsequent inferences. Because the source of her warrant for these beliefs is her color-vision itself, NSC maintains that they cannot by used in a Track-Record argument establishing that her color-vision is reliable. This claim, however, does not entail that her *perceptual beliefs* are unwarranted, or that their warrant isn't immediate. Hence, the dogmatist is free to endorse my NSC-based response to the bootstrapping argument regardless of whether one understands Roxanne's final inference as abductive or inductive.

iii. Two Worries about NSC

In this section I will consider two concerns the reader may have about NSC. The first is an argument due to Markie (2005) to the effect that NSC and principles like it are false. The second is the charge that my invocation of NSC in defense of dogmatism is *ad hoc.* I discuss each of these problems in turn.

It has been argued that principles like NSC are false. Markie (2005) proposes and rejects a principle that is similar to NSC:

> **IP:** Where a belief gains *prima facie* justification for S just from the fact that it was produced by a particular faculty (given e.g. the faculty's reliability, proper function, ability to provide the subject with evidencing experiences), the belief is not supporting evidence for S for beliefs concerning the reliability of that very faculty (Markie 2005, 414).

Markie presents two objections to IP. First, he claims that IP is false because we can and do "use our color perception to gain evidence that our color perception is reliable" (414). We need not consider whether or not Markie is correct on this point, since this objection does not implicate NSC. IP imposes a restriction on any means of coming to know that

one's faculties are reliable. NSC, on the other hand imposes restrictions on a *particular* method of testing one's faculties (namely by use of a Track-Record argument). Therefore, it is no objection to NSC to maintain that there is some way to come to know that one's color-vision is reliable that is in part supported by one's color-vision, so long as the particular way of relying on one's color-vision that NSC prohibits is admitted to be problematic.[45]

Markie's second objection does bear against NSC. Markie argues that IP conflicts with the following seemingly plausible claim:

> If it is reasonable for us to believe *p* and the truth of *p* increases the likelihood that another proposition, *q*, is true, then *p* is a reason (perhaps defeasible) for us to believe *q*. Just as money, however gained, still spends the same, so too reasonable beliefs, however gained, still epistemically support the same beliefs. (Ibid. 415)

In keeping with Markie's economic analogy, let us call this claim **Free-Spending**. As it turns out, White endorses something very close to this principle in support if Roxanne's abductive inference:

> Now if indeed I do know that my color-vision has been accurate on all of the many occasions I have used it in this test, then no matter how I may have come to know this, it surely supports the reliability of my color-vision. (White 2006, 546)

NSC is inconsistent with Free-Spending. NSC entails that, in the case of Track-Record arguments, one may have warranted beliefs (i.e. Roxanne's beliefs about the color of the individual cards) that cannot be used to support her belief that her color-vision is reliable. According to NSC, this can occur simply because these beliefs get their warrant from the

[45] Markie grants that IP could likely be reformulated to avoid this first objection (415).

very belief-forming method under investigation. If Free-Spending is true, then NSC must be false.

Markie and White do not provide an argument for Free-Spending and the principle has serious problems. Consider the following case:

> Joe wants to know about the current situation on the ground in Syria, and so consults the most recent Associated Press report. However, knowing that even a reliable source such as the AP can get things wrong from time to time, Joe checks a second source, a webpage that he has formerly found to be quite reliable, in order to verify the AP report. The webpage agrees with the AP report, and Joe takes this to provide confirmation that the AP report is accurate. Unbeknownst to Joe, however, the webpage's source for its information about the current conditions in Syria is the same AP report that Joe had previously read.

This case presents a counterexample to Free-Spending. Because Joe knows that the webpage is reliable, he is warranted in believing its reports. However, he cannot use those warranted beliefs to confirm the AP report's accuracy because the webpage's only source is the very report in question.[46]

It could be argued that Joe *can* reasonably confirm AP's accuracy in this manner because he is not aware of the connection between the webpage's reporting and AP's. Let us, therefore, consider what happens when Joe finds out about the connection. It seems that in this case he is no longer reasonable in utilizing the webpage's report as confirmation for what he learned from AP. However, since he knows that the webpage is reliable, it seems he is still warranted in believing its reports about the current conditions

[46] Wittgenstein mentions a similar sort of case: "As if someone were to buy several copies of the morning paper to assure himself that what it said was true" (*Philosophical Investigations I*, 265). Wittgenstein does not contend that newspaper is not a trustworthy source of information, or that one cannot use its testimony to form warranted beliefs. Rather, his point seems to be that whatever its epistemic merits, the newspaper cannot be used to confirm *its own* veridicality. This presents a similar problem for Free-Spending as I discuss here.

in Syria. Joe cannot use these warranted beliefs for certain purposes, such as confirming the AP's report or as further evidence for the general reliability of AP. However, it seems clear that if Joe's implicit assumption, that the webpage's information was independent of the AP report, were true, then Joe could have reasonably used the webpage as confirmation of AP's accuracy. What this shows is that, contrary to Free-Spending, in some cases the source of the warrant for one's beliefs can restrict the uses to which one can put them, in particular, one cannot use the testimony of a source to confirm the reliability of that very source.

Roxanne's case presents a similar sort of counterexample to Free-Spending. Consider a variation of the original case in which Roxanne already knows that her color-vision is reliable, and bases her belief in the color of the cards in part on this knowledge.[47] It should be clear that in this revised case Roxanne is warranted in believing that each card is the color it appears to be. However, it seems equally clear that Roxanne cannot deploy this warranted belief in her Track-Record argument to provide *additional* evidence of her color-vision's reliability.[48] Roxanne's background evidence about the reliability of her color-vision does not make bootstrapping any less problematic: the procedure intuitively does not provide Roxanne with any additional evidence for her color-vision's accuracy. However, if Roxanne were to support her confirmation clauses with an independent source, as NSC demands, then she could come to possess additional evidence that her color-vision is reliable via an NSC-respecting

[47] The reader may also add any stipulations that she deems necessary for Roxanne to be warranted in believing that a card is a given color.

[48] Weisberg (2010) uses similar points to attempt to establish that if bootstrapping is problem, it is a problem for *everyone*. See Cohen (2012) for a criticism of this argument.

Track-Record argument. Once again it seems clear that the difference between an acceptable Track-Record argument and bootstrapping is precisely what Free-Spending denies can be relevant, namely the source of the warrant for the confirmation clauses.

The general intuition behind NSC is clear; one cannot simply rely on the testimony of a source of information in order to conclude that that very source is reliable. However, I anticipate the criticism that NSC, as stated above, is *ad hoc*. NSC, it may be thought, is explicitly stated in order to allow the dogmatist a way to avoid sanctioning bootstrapping. I readily grant that this is so. It may also be claimed that NSC is an extraordinarily limited principle, in that it only applies to Track-Record arguments, and does not impose interesting general constraints on non-deductive arguments. This too, I accept. Thus, one might think that an appeal to NSC on behalf of dogmatism must be *ad hoc* and thus cannot be dialectically effective. This I deny. To accuse my response of being *ad hoc* is to misunderstand the character of my defense of dogmatism.

In stating and invoking NSC, I do not claim to have uncovered some deep truth or general constraint concerning abductive or inductive arguments. I do not claim that NSC is an independently interesting principle worthy of further investigation. I regard NSC as simply stating, in the most straightforward manner possible, what is meant to be problematic with bootstrapping. The point that I wish emphasize is simple: dogmatism is consistent with NSC. If dogmatism is consistent with NSC, then the view is not committed to sanctioning bootstrapping. Stated this way my response to the bootstrapping argument is deflationary. Dogmatism makes claims only about the conditions under which one can come to possess perceptual warrant: it does not make any claims concerning the norms governing warrant-preserving transitions amongst beliefs.

All the dogmatist need do in order to avoid sanctioning bootstrapping is to find some plausible diagnosis of the failure of Roxanne's final non-deductive inference that is consistent with the position. Since the dogmatist makes no explicit claims about the conditions for warrant-preserving inferences, it seems that virtually any plausible diagnosis will do.

The dialectical situation we are left with is as follows: critics have argued that dogmatism is committed to sanctioning an elicit form of reasoning: bootstrapping. If Roxanne's case is to establish this conclusion, then the objector must present some persuasive reason why the dogmatist is committed to endorsing Roxanne's final inductive inference from No Errors to the conclusion that her color-vision is reliable. No such argument has been offered. Thus, the dogmatist is free to endorse NSC as a means to avoid sanctioning bootstrapping.

iv. Another Problem for Roxanne

One might think that in focusing on Roxanne's inductive inference, I have ignored the true problem for dogmatism that the bootstrapping argument poses.[49] Cohen (2010) argues that the problem cannot simply lie in Roxanne's final non-deductive inference. Cohen's discussion suggests that Roxanne has already gone wrong if she has gotten as far as the premise that her color-vision was accurate on a particular occasion. This belief is arrived at solely by deduction from Roxanne's perceptually warranted beliefs in

[49] Some have attempted to answer the bootstrapping objection by arguing that Roxanne's reasoning goes astray at these early stages. Altschul (2011) defends such a view. His defense involves a commitment to a Humean sort of skepticism that I wish to avoid.

conjunction with her introspective beliefs. If the dogmatist accepts these intuitions and denies that Roxanne possesses warranted beliefs at these earlier stages, it seems as if the dogmatist is committed to denying CW. I have already acknowledged, this would constitute a significant theoretic cost.

In his 2010 paper, Cohen's presentation of the bootstrapping procedure makes use of different intermediate premises than the version of the argument already considered. Cohen's subject, (call him Stew), reasons as follows:

> (1*) Card 1 looks red.
> (2*) Card 1 is red.
> (3*) Card 1 looks red and is red.
> (4*) My color-vision worked correctly (on this occasion). (Cohen 2010, 142)

Stew's (1*) and (2*) are based in an analogous way to Roxanne's (1) and (2) respectively. The inference from (3*) to (4*) is purported to be deductive. After a suitable number of trials the subject infers:

> **No Errors*:** My color-vision worked correctly every time, i.e. I made no errors. (Ibid)

Stew then concludes on the basis of No Errors* that his color-vision is reliable. Cohen, in considering a response to the bootstrapping argument like mine that focuses on Roxanne's final inference, suggests that it is implausible that Stew could even get as far as (4*) while retaining warrant. (Ibid. 144) I agree. However, this presents no problem for dogmatism because on any reasonable interpretation of "worked correctly" (3*) simply does not entail (4*). Stew's color-vision could produce a representation of red, the card could be red, and yet his color-vision could have failed to work correctly.

Given Cohen's statement of No Errors*, it seems clear that by claiming that his vision "worked correctly" in a particular case, Stew means that his color-vision has not made an error. This conclusion is not entailed by the fact that the color of the card is the same color as it visually appears to Stew to be. Consider an analogous point with regards to belief. A belief can be true without it being the case that the believer has arrived at it in an error-free manner; the mere truth of my belief entails nothing about how I came to believe it. In good cases, I have arrived at the belief without making errors, but in bad cases I may have made errors that (as it turns out) still lead me to a true belief. Suppose that I come to believe *p* by repeatedly affirming the consequent, or by making use of other logically fallacious inferences. However, suppose that, for reasons I have not yet considered, *p* is true. In such a case I have formed a belief through a series errors, and yet my belief is still true.

Similarly, one's perceptual experience as of a red card can be veridical, even if one's perceptual system has functioned improperly, and/or made errors in processing the information received by the proximal stimulation of the retinas. Suppose that I look at a red card and my retinas are stimulated in the way that they normally are when there is a red object in front of me. However, in processing that information, due to numerous strong localized electromagnetic fields located near my optic nerve and V1 cortex there are a series of malfunctions in my ordinary visual processing. The processes still produce a visual experience as of a card of a certain color, but the process is ridden with malfunctions. By sheer luck, however, all of these errors compound in such a way that the resulting experience is as of a red card. Despite the fact that my visual system has massively malfunctioned, it has managed to produce a veridical perceptual experience.

Thus, the mere fact that experience matches the actual color of the card does not imply that my visual system has not made an error.

Perhaps a charitable interpretation of Cohen's objection would rephrase his claims in terms of *accuracy* rather than *error*. Put this way, the objection would be that it is problematic for Roxanne to reason from (1*)-(3*) to the conclusion:

(4**) My color-vision was accurate on this occasion.

There is some ambiguity in the word "accurate" that we must contend with. Consider a marksman firing at a target: we can be concerned with the accuracy of a particular shot (i.e. it hit the target) or the accuracy of the gun or the scope (i.e. it will reliably put a bullet where aimed in ordinary conditions).[50] In the case of a belief-forming method such as color-vision, one could evaluate it's accuracy in similarly diverse ways: (1) one could wonder whether or not the belief it produced on a particular trial was true, or (2) one could wonder whether or not, in producing a given true belief the method manifested some general competence or reliability in securing true beliefs. To dispense with this ambiguity, it is useful to introduce a useful distinction provided by Ernest Sosa (2010). Sosa reserves the term *accuracy* for an evaluation of the first type (i.e. whether the belief is true), and the term *aptness* for the second (i.e. whether the truth of the belief manifests some competence). On neither reading of (4**) is there a significant problem for dogmatism.

If (4**) is read as attributing *aptness* to one's color-vision, then it is not entailed by (3*). Reading (4**) in this stronger way certainly does seem to make it implausible that Roxanne could come to be warranted in believing it in the manner described.

[50] This analogy, like the distinction, is due to Sosa (2010).

However, so understood (4**) is not entailed by the simple fact that her color experience matched the color of the card. One cannot deduce the aptness of Roxanne's color-vision from the truth of (3*) for the same reason one cannot deduce (4*)—and its claim that Stew's color-vision "worked correctly" in a given trial— from (3*). Roxanne's color experience matching the card does not entail anything about the competence or reliability of her visual system. Someone could have a visual experience as of a red card when there is in fact a red card there, but that experience could be the product of a highly incompetent, unreliable visual system with regards to colors.

Consider a case in which a person's visual system delivers color experiences at random when presented with colored objects. There is nothing reliable or competent about this subject's color-identifying abilities. This is compatible with lucky matches between her randomly generated color experiences and the actual color of the objects viewed in particular cases. The mere fact that a belief that results from the exercise of one's color-vision is true does not entail that one's perceptual system manifests *aptness* in Sosa's sense. At best it seems as if Cohen could claim that there is a strong non-deductive inference that the dogmatist must sanction for this stronger accuracy claim, but this too seems dubious. The same intuition that underlies NSC could be deployed in the present context to rule out such inferences. If it is objectionable to rely solely on a source to confirm its own reliability it seems equally problematic to rely solely on a source to confirm its own aptness.

If we understand (4**) as making the weaker claim to *accuracy* in Sosa's sense, it becomes unclear what the problem is meant to be. (1*) is stipulated in the construction of the case to be warranted. The dogmatist claims that (2*) is warranted in virtue of

Roxanne's undefeated visual experience. Given these stipulated features of the case, it is hard to see how there could be anything problematic about Stew's belief in (4**) understood as a claim about accuracy in Sosa's sense. According to the weaker reading of (4**), for one's color-vision to be accurate it simply needs to be the case that one's color-vision delivered a true belief in a given case. To implicate dogmatism, there must be some clear problem with Stew's conclusion that his belief in the color of the card in this trial is true. But, if Stew is warranted in believing that the card is red, he is also warranted in believing it to be *true* that the card is red. So long as one possesses the requisite concepts, it is unclear how one could come to be warranted in believing *p* without simultaneously being warranted in believing that *p is true*. However, on the accuracy reading of (4**), what the claim amounts to is that the card looks red and it is true that the card is red. If Stew is warranted in believing the former, it seems as if he is unobjectionably warranted in believing the latter.

Any residual appearance of a problem may be the result of one of two factors: the first is that one is tempted to conflate *accuracy* with *aptness*. In virtue of being warranted in believing *p*, Roxanne is also warranted in believing that her belief-forming method was *accurate*. While I have argued that it is not objectionable for Roxanne to conclude that her color-vision was accurate in a given trial, it does seem problematic for her to infer that it was *apt*; aptness implies competence or reliability that the mere truth of the belief resulting from her color-vision does not support. But the claim that her color-vision is apt is not entailed by (3*), and so the dogmatist is not committed by CW to accepting that Roxanne is warranted in believing the stronger version of (4**).

The other factor has to do with the potential use of Roxanne's belief in (4**) in further reasoning about the reliability of her color-vision. One might think that if she can know that her color-vision was accurate in the way that I have discussed then she will have no trouble producing a strong case for the reliability of her color-vision on the basis of these beliefs. But, as I argued in the previous section, this appearance is illusory. Reasonable constraints on the manner in which Track-Record arguments can proceed rule out the possibility that Roxanne will be able to utilize this procedure to warrantedly conclude that her color-vision is reliable.

8. Conclusion

I have considered six prominent objections that have been raised against dogmatism and found all of them wanting. As such we have encountered no reason to suppose that the claim that having a perceptual experience as of p provides one with *prima facie,* defeasible warrant to believe p is itself seriously problematic. This chapter has been devoted to defending the core dogmatist thesis from direct attacks. However, there are many different ways of developing this claim into a more complete epistemological theory of perception. Even if there is nothing problematic about this core dogmatic claim, it is possible that there will be significant problems attaching to the various attempts to answer what I take to be key questions about the dogmatist approach to perceptual epistemology. In particular, I have thus far been entirely silent on a very important issue: in virtue of what does one enjoy the perceptual warrant that the dogmatist claims one does? Do all perceivers without fail enjoy this warrant, or must one satisfy some additional conditions in order to possess the warrant in question? Different

answers to these questions may well involve commitments that make the resulting view unpalatable. I turn now to a discussion of a variety of dogmatist theories and the comparative merits of each.

Chapter III. Internalism and Dogmatism

1. Meta-Epistemology and Dogmatism

Thus far my discussion has focused on explicating and defending the idea that perceptual warrant is immediate and defeasible. I have defended a thesis about the conditions according to which a perceiver can form a warranted belief on the basis of their perceptual experiences; the conditions being that they have a perceptual experience with the right content, and no defeating evidence that undermines or overrides the warrant the experience provides. I believe that this account is correct, but as epistemologists we should not be satisfied with a merely extensionally correct account of when perceivers gain warrant from their perceptual experiences; we should also seek to explicate the underlying reasons that this warrant exists. Supposing, then, that the dogmatist is correct in claiming that perceptual experiences warrant beliefs in their contents, the question we must now ask is *why* this is the case. What is it about perceptual experiences that ground the sort of warrant that I have argued that they provide? In virtue of what do perceptual experiences warrant basic perceptual beliefs? Call this, the **meta-epistemological question**. The terminology is meant to emphasize that this is not a question about the conditions a person must consciously and reflectively meet in order to be perceptually warranted in believing something. Those conditions are adequately captured by the formulation of dogmatism I have defended throughout the preceding chapters. Rather, this is a philosophical investigation into the nature and origins of the warrant supplied by perception.

The meta-epistemological question must be confronted by any epistemological theory, but one might think it is especially pressing for theories that advocate immediate

warrants as dogmatism does. It seems that we have a working understanding of how one belief can support another via its logical and evidential relations to other beliefs. But it can seem utterly mysterious how a belief can come to be warranted when it is not based on another belief. Does the warrant spring forth fully-grown from some normative void? What is the origin of this seemingly magical epistemic value? Given that the dogmatist is committed to the claim that no belief is needed to provide epistemic support for basic perceptual beliefs, what is the explanation for the warrant she claims that such beliefs possess?

It may be thought that this perceived special difficulty is due to a simple confusion. According the dogmatist, perceptual beliefs are not without any epistemic basis: they are based on *perceptual experiences*. Perceptual experiences, though not beliefs, are contentful mental states capable of entering into evidential relations with basic beliefs in much the same way that beliefs can be used to support other beliefs. While this is an important point and should not be overlooked, it doesn't remove the teeth of the meta-epistemological challenge to the dogmatist; it simply pushes the problem back a step. The challenge is renewed if one wonders: in virtue of what are perceptual experiences capable of warranting beliefs? The answer cannot just be that the experience and the belief share the same (or approximately the same) contents. There are many contentful non-belief states: desires, imaginings, and idle thoughts, to name a few. Presumably, if I were to close my eyes and wish that there were a cup on the table, I could not form a warranted belief that there is a cup on the table on that basis. If I were to idly consider the proposition that Obama is on a water-slide (as I just did) I am not, thereby warranted in believing that Obama is on a water-slide. The meta-epistemological

question can thus be presented as a challenge for the dogmatist: in virtue of what do perceptual experiences provide warrant for believing their contents?

The meta-epistemological question is of independent interest, but it is also the primary basis for determining which, of the many dogmatic theories currently on offer, is the most defensible. "Dogmatism," as I have defined the term, casts a wide shadow. In line with Pryor (2013), I have used the term so that it applies to any theory that allows for immediate, defeasible perceptual warrant. Under this umbrella are an array of views, including (but not limited to): internalist positions such the phenomenal conservatism of Michael Huemer (2001), the phenomenal approach of Declan Smithies (2014), externalist theories such as Alvin Plantinga's proper function theory (1993), as well as the approach favored by Tyler Burge (2003) and Peter Graham (2010) based on the constitutive and causal connections between the formation of the contents of perceptual experiences and the evolutionary environment of the organism. The task of this chapter and the next is an investigation of which of these many positions has the best chance of providing a satisfactory answer to the meta-epistemological question. The current chapter will focus on internalist approaches to the problem and argue that the introspectively accessible elements of perceptual experience are not sufficient to provide a satisfactory answer to the meta-epistemological question. The next chapter will develop an externalist account of perceptual warrant that I call **competent dogmatism**.

Particular answers to the meta-epistemological question can have various undesirable consequences. In many cases, answering the meta-epistemological question can commit one to a view that is extensionally incorrect. One's attempt to explain perceptual warrant could result in a view that is *overly permissive* or *overly strict*. In the

former case, one's view will allow that intuitively unwarranted perceptual beliefs are warranted, while in the latter one's view entails that certain perceptual beliefs that seem perfectly acceptable are unwarranted. On the other hand, certain answers to the meta-epistemological question fail on their own terms if they do not plausibly explain the existence of perceptual warrant. In this chapter, I will argue that internalist versions of dogmatism fail in one or both of these ways.

2. Introspection and the Internalist Intuition

The term "internalism" has been used in epistemology to refer to many different varieties of views,[51] but in the present context I mean to discuss views that maintain that what warrants one's beliefs must be either introspectively accessible or reflectively available to the subject. For my purposes, then, let us define internalism as follows:

> **Epistemic Internalism:** What determines the epistemic status of S's beliefs is introspectively accessible (at least in principle) to S. [52]

[51] For a useful survey of these positions, see Conee and Feldman (2001) and (2004) Ch. 3. Connee and Feldman defend a view called *mentalism* according to which what determines one's epistemic status are solely determined by one's mental states some relevant features of which may not be introspectively accessible. The variety of internalism which is my focus here is what they call *accessibilism* according to which the epistemic difference makers must be introspectively accessible to a subject as I define it above. The reason for this focus is that the motivations appealed to by internalist dogmatists are typically at odds with a mentalist approach. Both the New Evil Demon problem and Huemer's internalist argument can be marshaled against a view that allows for mental but introspectively inaccessible differences to matter epistemically.

[52] Epistemic internalism is to be distinguished from another form of internalism: *content internalism*. The latter is a thesis concerning constitutive conditions of possessing representational states of a certain kind. Very roughly, content externalism maintains that representational contents depend on constitutive causal relations to extra-mental objects. Epistemic internalism is, strictly speaking, consistent with content externalism. This is because the contents of our mental states are typically introspectively accessible to agents of sufficient cognitive complexity, even if they are

Let a proposed condition on warranted belief be *externalist* if and only if it need not be introspectively accessible to the believer in order for her to be warranted; knowing that an externalist condition on warranted belief is met could require empirical investigation on the part of the subject; (e.g. whether her visual experiences are reliably veridical). There are powerful intuitions that lend support to internalism so construed. Take two oft-cited externalist conditions: that one's belief-forming methods are reliable, or that one's beliefs be caused in an appropriate way. Even if my actual experiences are reliable and/or caused in the appropriate way, everything could seem the same to me introspectively, while I was radically deceived and/or my experiences were formed based on a deviant causal chain. Any position which maintains that proper causation and/or reliability are necessary condition on warrant must, therefore, concede that introspectively identical subjects can differ in the warrant they possess for their beliefs. This strikes many as untenable.

To illustrate the problem: consider my brain-in-a-vat duplicate, BIV-Tim. BIV-Tim has spent his entire existence being fed experiences by electrodes hooked up to his brain that exactly mirror my own.[53] Unfortunately most of BIV-Tim's experiences are non-veridical and/or caused inappropriately. First, what ought BIV-Tim (epistemically speaking) believe? If, as is plausible, one ought to believe all and only that which one is

---

partly individuated externalistically (Burge 1988), (Brueckner 2010, selections 15-17 and 23).

[53] Such a case may be inconsistent with content externalism, a thesis that I endorse. However, one could revise the case to consider a brain that was envatted as a young child after the contents of her perceptual states had been fixed by causal interaction with ordinary physical objects. I ignore this complication in the present discussion, though I return to it briefly in Chapter IV Section 8.

warranted in believing, it seems as if there is no good externalist answer to this question. A reliabilist or causal-historical externalist cannot say that BIV-Tim ought to believe what I do; according to her, my beliefs are warranted in part based on their etiology and/or the reliability of my faculties. But while *my* beliefs are formed by reliable processes and caused appropriately, BIV-Tim's are not. According to such accounts, then, I am warranted, but BIV-Tim is not. But it seems callous to claim that BIV-Tim ought not believe *anything* about the world. After all, things seem to him just as they do to me. If he believes what I do, he seems to be reacting in an entirely appropriate way to his experiences. It is not his fault that, unbeknownst to him, his experiences are radically deceptive. He has no means to find this out and no means to change it. Given his experiences, BIV-Tim ought to believe just what I do. If this is correct, the externalist is faced with the unenviable task of explaining why BIV-Tim ought to form unwarranted beliefs.

Huemer (2006) does an admirable job of isolating the intuition at play in this case. He constructs a schema for developing (purported) counterexamples to any externalist condition on epistemic warrant. Let E be an externalist condition on warranted belief. Let C be an enumeration of the sufficient conditions for warrant including this condition. If E is a genuine condition for warranted belief, then the following case seems possible to construct:

1. S satisfies E with respect to p, but not to q
2. S satisfies all other parts of C with respect to both p and q.
3. S satisfies no other sufficient conditions for warranted belief in p or q.
4. It seems to S as if he is in the same epistemic position with respect to p as he is to q, and he has no reason to suspect that one is more warranted or likely to be true than the other.
5. S ought (epistemically speaking) to believe p and not believe q (Huemer 2006, 149).

Externalists face a dilemma regarding any case satisfying this schema. On the first horn, warrant is radically decoupled from what one ought to believe. In such a case, S is warranted in believing p and not warranted in believing q, but ought to believe both (or neither). But this seems to be a sort of conceptual confusion; warrant is a normative concept and one might think that it makes little sense to claim that a subject both ought to believe something and fail to be warranted.[54] On the second horn the externalist must say that S ought to believe p and ought not believe q even though both propositions seem, from an internal perspective, epistemically on par and just as likely to be true to S. But it seems that, given S's epistemic state he is in no position whatsoever to identify the purported difference in warrant, and so S has no rational grounds on which to judge p more likely to be true than q. Since p and q are, from S's perspective just as likely to be true, it seems irrational for him to take different epistemic attitudes towards them. Since externalism fails on both horns, one ought to adopt an internalist view of warrant.

The preceding is a general argument for epistemic internalism, and seems to apply equally well to the special case of perception. The following case is an instance of Huemer's general schema. Smith, an entirely normal adult human who possesses all the ordinary beliefs of an ordinary adult human, has a normally caused and otherwise entirely normal visual experience as of a red ball on the table, no defeating evidence for his belief, and there is a red ball there. Intuitively, Smith is warranted in believing that there is a red sphere on the table (I know of no view, excepting radical skepticism, according to which Smith is not warranted in his belief). Now suppose Smith blinks, and in that instant two

[54] I return to this point in the concluding section of this chapter. With a bit of clarification, it can be shown that there is no serious problem in adopting this suggestion.

things happen. The Flash speeds by and "borrows" the ball to aid in his crime fighting efforts, and (in an entirely unrelated coincidence) Martians send a beam into Smith's brain stimulating the occipital lobe such that he will, on opening his eyes, have an experience as of a red ball completely indistinguishable from the experience he had a moment before. From Smith's perspective his belief that there is now a red sphere on the table is on par with the similar belief he had before blinking. He cannot identify any difference between them that would cause him to suspect that one belief was more likely to be true than the other. However, given that his experience as of a red sphere now is aberrantly caused, according to externalist theories that mandate proper causation for warranted belief, Smith fails to be warranted. But again, it seems implausible to claim that he ought not believe that there is a red sphere on the table or, equally problematically, that he ought to form an unwarranted belief. This case is meant to be an illustration of the application of Huemer's schema to the case of perception, and for that purpose I have chosen a causal criterion like that advocated by Goldman (1967) as an example. However, it should be clear that a similar sort of case could be generated for any externalist condition on warrant. These considerations are sufficient to provide an internalist version of dogmatism with at least some initial motivation in contrast to its externalist counterparts.

Let us turn, then, to the meta-epistemological question and the prospects for an internalist answer to it. An internalist answer to the meta-epistemological question may only utilize resources that are internally accessible to the agent. Suppose that there is some state of affairs C that partly explains S's warrant for believing p and q, and that C is not introspectively accessible to S. Using Huemer's schema it is a trivial matter to

construct a counterexample to such an account. Since S does not know whether C obtains, it is possible that, for all S can be aware of, C obtains with regards to p, but not with regards to q. Since S's warrant for believing p and q purportedly depend on C, then S will not be warranted in believing p but not in q. But this is just the result that internalism was meant to avoid. Thus, whatever explanation the internalist seeks to give in answer to the meta-epistemological question must concern only that which is accessible to the agent.

A second constraint on internalist attempts to address the meta-epistemological question is entailed by the immediacy of the perceptual warrant that is the core of the dogmatist thesis. Some care must be taken in formulating this restriction. Being a dogmatist, an internalist of the kind that interests us here cannot make use of other mental states the agent has in order to answer the meta-epistemological question *if her perceptual beliefs depend for their warrant on the warrant of those other states*. This does not explicitly rule out background beliefs and other mental states playing some role in explaining perceptual warrant so long as that role does not depend on the warrant those states themselves enjoy. Most theorists have taken this restriction to mean that other mental states of the subject (e.g. background beliefs, memories, intuitions, etc.) cannot play a role in explaining perceptual warrant. This assumption is reasonable; it seems likely that if other introspectively accessible mental states are to play a role in explaining perceptual warrant, they must themselves be warranted. It is unclear how entirely unwarranted memories, intuitions or beliefs could help to explain why we are entitled to trust our perceptual experiences. But if the warrants for these states play a role in explaining warrant for perceptual beliefs, the warrant is no longer immediate.

Dogmatists of internalist leanings, therefore, typically attempt to answer the meta-epistemological question by appeal to some introspectively accessible features of the perceptual experiences themselves. The resources here are sparse. The reliability veridicality of one's perceptual systems, the causal history of the experience (e.g. that it was caused by a physical object, that it is the result of a properly functioning visual system, etc.) and so on are ruled out. One could find oneself in a state of affairs introspectively indistinguishable from a normal perceptual experience in which these conditions are not met. Nothing about the regular connection of the experience to the world, nor the experience's role in the functioning of one's cognitive system will be introspectively available. Indeed, even its status as a *perceptual* experience is not introspectively available in the required sense. It is possible to be experiencing a realistic hallucination or to dream in a manner that is introspectively indistinguishable from a perception. Hallucinations and dream-experiences are not perceptual experiences, and so there will be cases in which one cannot know by introspection alone whether or not one is perceiving, or merely seeming to perceive some event.

Michael Huemer (2001, 66) distinguishes three introspectively accessible elements of a perceptual experience:

1. Representational Content
2. Qualitative Character
3. Forcefulness

The question that I address in the remainder of this chapter is whether any of these can serve to provide a plausible explanation for the warrant that perceptual experiences supply for basic beliefs. I have already argued above that the representational content

alone cannot explain the warrant provided by perceptual experiences.[55] Therefore, in order to answer the meta-epistemological question, the internalist must appeal either to its phenomenal character, or its forcefulness. Taking the phenomenology or qualitative character of experiences first; there seem to be two possible positions available to those that think that the phenomenology of one's perceptual experiences explain one's warrant to believe their contents: first, one could maintain that there is something epistemically important about the individual phenomenal character of particular experiences; e.g. there is something epistemically important about *what it is like to see green* such that it entitles one to believe that there is a green thing there. In contrast, one could maintain that there is a phenomenal component to perceptual experiences generally, a "truthy feel," that is capable of explaining the warrant they provide beliefs. According to this second option, the strategy is to assimilate what Huemer calls *forcefulness* into the phenomenal character of one's experiences. Huemer and others phenomenal conservatives, however, typically maintain that forcefulness is a non-phenomenal, feature of experiences and other mental states that pick out *sui generis* a category of propositional attitudes he calls "seemings." Such states are introspectively accessible and purportedly capable of entitling beliefs.

In the Section 3, I argue that views that the attempts to explain perceptual warrant in terms of the particular phenomenal character of perceptual experiences fail to provide an adequate account of the nature of perceptual warrant. Such views fail on explanatory grounds because they have no resources with which to explain how subjects whose perceptions have radically different phenomenal characters can both be warranted in the same beliefs. In order to illustrate this problem I construct a hypothetical case on the

[55] Though, as I have previously stated and will emphasize in the present context, it does play a role in fixing *which* beliefs a perceptual experience warrants.

basis of the reported experiences of blindsighted subjects that indicates the possibility of a *single* subject having experiences with radically different phenomenology character in different parts of her visual field. I argue, further, that no appeal to *representationalism*, (a view in the philosophy of mind which maintains that phenomenology is some kind, or bears some close relationship to representation content), can aid such a view.

Section 4 confronts Michael Huemer's theory of phenomenal conservatism, according to which it is the *forcefulness* of one's experiences that explains their warranting power. I argue on the basis of cases presented by Markie and Siegel that phenomenal conservatism fails extensionally in that it is to permissive in ascribing warrant. I consider and reject a response to these cases advocated by both Huemer and Tucker.

3. The Phenomenal Approach

i. The Phenomenal Approach and Phenomenal Contents

One of the most introspectively striking features of our perceptual experiences is their phenomenology. The *feel* of heat, the *look* of deep blue, the overpowering *stench* of a skunk nearby, and so on. These experiences are deeply tied up with the manner in which we interact with and experience the world. It is natural, therefore, to think that phenomenology must play a central role in any epistemic account of perceptual warrant. Furthermore, the introspective ubiquity and salience of phenomenology in perception seems tailor-made for an internalist answer to the meta-epistemological question for dogmatism: perceptual experiences warrant basic perceptual beliefs in virtue of their phenomenology.

We must be careful, however, in framing our discussion of this suggestion. In the first place, we should distinguish between the *representational content* and *phenomenal character* of experiences. The representational content of an experience is the way the experience represents the world to be. Typically when I look at a red ball, my perceptual experience is as of a red sphere in front of me. Sometimes referred to as the intentionality of experiences, the representational content determines what the experience is about. An important feature of such contents is that they are individuated, in part, by their veridicality conditions. My experience is as of a red sphere just in case, were there to be a red sphere before me, my experience would be veridical.

There are many ways to characterize the phenomenal or qualitative character of experience, though any way of doing so besides simply gesturing in its general direction is bound to be controversial.[56] Therefore, what I say here is meant to do nothing more than isolate, mostly by ostension, the element of experience that will be our subject. The phenomenal character of an experience is individuated by *what it is like* to have the experience. It is the element of an experience one cannot appreciate or fully understand until one actually *has* the experience in question. For instance, if one has never smelled a skunk (and has never had an experience of something with a comparable smell), there is very little that anyone can do to explain exactly how it smells. One can point to its deeply unpleasant nature, how, if one is too close overpowers all other attempts at olfaction, and how it seems to permeate the environment, but it seems that no explanation along these lines can do justice to the *particular* character of a skunk's stench. Until one

[56] We will have some cause to dabble in some of this controversy, if for no other reason than to do justice to the epistemological views under consideration which depend crucially on substantive claims about the nature of phenomenology, and in particular its relation to representational content.

smells it for oneself, one will simply be in the dark about how a skunk smells. Similar things may be said about other sorts of experiences: seeing a color, hearing a trumpet blast, hitting one's funny bone, having an orgasm, and so on. This element that must be experienced in order to be fully understood is what we will mean here by the *phenomenal character* of an experience.

There are ways in which experiences can share contents, but differ phenomenally. The least controversial way to illustrate this is to consider cross-modal perceptual contents. I can come to believe on the basis of a particular visual experience that something round is on the table in front of me, and I can also do so, (if my eyes are closed, or I am blind) by reaching my hand out and feeling it. We can say of both the tactile and visual experience that it is an experience as of something round, however the particular phenomenal character of each differs radically, so much so that cross-modal comparisons of phenomenology often seem impossible. In such cases, even though the *phenomenal character* of one's experiences is very different, the *representational content* is the same (at least in part). [57]

With this distinction in mind, I can now state the view that will be the subject of this section's investigation:

> **The Phenomenal Approach (PA):** The immediate warrant furnished by a perceptual experience as of p for a belief that p is explained by the particular phenomenal character of the experience that p.

[57] What I say here is meant to be neutral in the debate over the merits of *representationalism*. As I point out in what follows, representationalists can plausibly accommodate cross-modal representational contents by pointing out the various representational differences that any two experiences from different sense are bound to have. Thus, they can accommodate the phenomenal difference by positing *other* representational differences, besides the one's that they share. On this view, the phenomenology of an experience is a holistic matter, determined by a complex interaction amongst the various contents of one's experience.

According to this answer to the meta-epistemological question it is the phenomenal character of individual experiences: the *look* of red or the *feel* of heat that explains the epistemic support provided by perceptual experiences for the beliefs that some object is red or there is something hot nearby. Furthermore, according to this proposal it is not enough that the subject have some phenomenal experience or other, but rather that the *particular* phenomenal character of an experience plays a significant role in providing warrant for the resulting belief. I will argue that PA and more plausible successor of the view fail as explanations of our basic perceptual warrant.

ii. The Inverted Spectrum

One can begin to make trouble for PA by considering a considering a classic philosophical thought experiment:

> **Inverted Spectrum**
>
> When Andy looks at a stop sign, he has an experience with the same phenomenal character Ollie has when he looks at fairways at Pebble Beach golf course. Both of them are competent, native, English speakers. When Andy and Ollie look at a stop sign they agree that it is "red." And when they look at healthy grass they agree that it is "green." Andy and Ollie are inclined to make the same or similar judgments in many other cases of similarly colored objects. They will make largely the same comparative judgments (the stop sign is the same color as this ball), and they both share substantive agreement about the color of objects of this variety with other members of their language community.

Allow me some observations about the case. Andy and Ollie both believe that the stop sign is red and that the grass is green. Neither Andy nor Ollie are wrong about the color of these objects and they both have true beliefs about the color of the grass and the sign.

They share the same concepts with regards to the color of the objects; they are able to converse intelligibly about the color, come to an agreement about the color and report their agreement to others in their language community. Their perceptual experiences of the stop sign and the grass are each *veridical*. That is, when Andy looks at the sign his experience is veridical just in case the sign is red, and when Ollie looks at the sign, his experience is also veridical just in case the sign is red. Since the sign is red, both of them are enjoying a veridical color-experience. There is something about Andy's and Ollie's experiences that they share, and there is something that they do not. What is alike in their visual experience is their representational character. On the other hand, their perceptual experiences differ in their phenomenology. There is a difference in what it is like for Andy and Ollie when they look at a stop sign even though their experiences are both veridical in just the same conditions.[58]

Many are inclined to deny that Inverted Spectrum is possible. Typically this is due to some thesis about the nature of minds that entails the impossibility of two subjects differing mentally while remaining the same in some other way (e.g. physically, behaviorally, functionally, and so on). Such grounds for denying the possibility of Inverted Spectrum are irrelevant here. The case as I have stated it allows for a host of differences between the subjects. In particular, the case as stated does not assume that the subjects are behavioral, physical, neurophysical, or functional duplicates. These differences can be allowed for, so far as I am concerned here, and can even help explain Andy and Ollie's phenomenal differences. I am interested in the epistemological rather than the metaphysical ramifications of the case. As such, many of the grounds usually

[58] I am aware that some of these claims about the case are controversial. I will consider the relevant controversies in due course.

adduced to argue that cases like these are impossible may be safely ignored as long as we make the quite plausible assumption that physically and/or functionally different beings can have perceptual states and beliefs with the same color contents.

The *prima facie* problem for PA that Inverted Spectrum raises is easy to state. Suppose that Andy is looking at grass and Ollie is looking at a stop sign. According to my description of the case they have phenomenally identical experiences with regards to color, but despite this, the experiences warrant different color-beliefs: Andy is warranted in believing that what he is looking at is green and Ollie is warranted in believing that what he is looking at is red. Any account that purports to explain the warrant Andy and Ollie have for their beliefs must, at minimum, explain why their experiences warrant different beliefs. However, since their experiences are phenomenally identical, one cannot appeal to it in order to explain the epistemic differences between the subjects. Similarly, suppose Ollie turns from his view of the stop sign to look at the healthy grass of his lawn. Now *both* Andy and Ollie are warranted in believing that what they are looking at is green, but they are so warranted on the basis of experiences with different phenomenal characters. Once again, it is mysterious how PA can accommodate this result. Any explanation of the warranting power of Andy's and Ollie's perceptual experiences should be capable of explaining how their experiences succeed in warranting the same belief. However, one cannot appeal to the phenomenal character of their experiences to ground their epistemic similarity precisely because the phenomenal character of their experiences differs.

One diagnosis of the failure of PA in this case is that *which* beliefs a perceptual experience warrants depends on the representational content of the perceptual experience.

Thus, the plausibility of PA as an answer to the meta-epistemological question depends crucially on one's view about the nature of perceptual phenomenology, in particular the degree to which one takes the phenomenal character to determine, be determined by, (or even be identical to) the representational content of the experience. In the parlance of professional philosophers, PA becomes more plausible if one adopts some version of *representationalism* or *intentionalism* about the nature of phenomenology. According to such a view phenomenology is not simply *raw feels* or *qualia*, as perhaps the statement of the inverted spectrum case above assumes, it has genuine representational content that is closely related to its particular phenomenal character. According to the representationalist, two experiences cannot be phenomenally the same without also having the same representational contents. David Chalmers (2006, 50) introduces the following useful terminology to capture this purported relation between phenomenology and content:

> **Phenomenal Content:** For any experience *e* and content *c*, *c* is a *phenomenal content* of *e* just in case every experience with the same phenomenal character as *e* has *c*.

In other words, having a certain phenomenal character entails having a given representational content. Let us call a view *representationalist* just in case, at minimum, it endorses the existence of phenomenal contents. This characterization is useful starting point. That perceptual experiences have phenomenal contents is entailed by a views that advocate for a much stronger connection between phenomenology and content such as those advocated by Dretske (1995) and Tye (1992) according to which *all* phenomenology is simply a kind of representational content. Trivially, if the phenomenology of an experience *just is* a part of its content, then any experience with

phenomenal character will have some phenomenal content: namely, the content that is identical to the phenomenal character. However, one need not endorse this strong version of the view in order to think that there are such things as phenomenal contents. For instance, one could think that phenomenology is causally or constitutively related to certain content determining aspects of visual experience without thinking that one can simply identify the phenomenal with the representational.

With the notion of a phenomenal content in hand, we can amend the phenomenal approach as follows:

> **The Phenomenal Approach (PA*):** The immediate warrant furnished by a perceptual experience as of p for a belief that p is explained by the phenomenal contents of the experience that p.

The Inverted Spectrum argument presented above simply does not address PA*; the case was described in terms of the shared phenomenal character of Andy's and Ollie's experiences and PA* makes use of the quite distinct notion of *phenomenal contents*. Furthermore, it seems that there is no non-question-begging way in which to restate the argument so that it bears against PA*. Consider the case in which Andy is looking at the stop sign and Ollie is looking at healthy grass. As I described the case above, their experiences have the same phenomenal character: according to my description of the case, the representational contents of their experiences differ. However, a proponent of PA* will feel no compulsion to grant that this is possible. In order to explain the warrant that perceivers ordinarily have for their visually formed beliefs about color, PA* must insist that *color contents* of visual experiences are *phenomenal contents*. If color contents of visual experiences are phenomenal contents, then it will be impossible for two subjects to have an experience with the same phenomenal character while the color contents of

their experiences differ. Therefore, the case I have described will be deemed impossible by proponents of PA*. Moreover, this denial is not *ad hoc* as it is based in part on a live and independently supportable view about the nature of phenomenology. The conjunction of the phenomenal approach to and a mild version of representationalism is, therefore, immune to the argument based on Inverted Spectrum.

Since the response I have sketched depends on a version of representationalism, one way to renew the objection to PA* is to argue against representationalism itself. If representationalism is false, it is quite difficult to see why Inverted Spectrum as I have presented it is impossible given the wide range of differences physical and functional that the case allows for. However, this seems like a questionable strategy in the present context. Cases like Inverted Spectrum feature heavily in the debate over whether or not all phenomenal states are contentful, and whether or not perceptual experiences have phenomenal contents. As such, appealing to such cases in denying representationalism risks begging the question against proponents of PA*. Of course, there are other ways in which one could object to representationalism[59], but there is a voluminous literature concerning the soundness of such arguments that would take us quite far afield from our present epistemological project. An approach that avoids such difficulties is to be preferred in the present context.

[59] For my part, I am convinced by examples such as orgasms and headaches that the attempted reduction of some representationalists of phenomenology to representational contents is, at best, highly problematic. I am also convinced, for reasons that will hinted at in what follows, that phenomenal contents do not play an important or interesting role in many of the core questions in the philosophy of perception. It should be noted, however, that despite what I argue here I take an investigation into the nature of phenomenology, its relation to other features of the mental, and its relation to physical objects to be a difficult deeply interesting domain of inquiry in its own right. My point in the present context will be that appeals to phenomenology are not illuminating for the purposes of explaining perceptual warrants.

In what follows I will argue that the debate between representationalists and non-representationalists is orthogonal to the epistemological question of the role of phenomenology in warranting perceptual beliefs. No matter how close one chooses to make the connection between contend and phenomenology (including the limiting case in which phenomenology is identical to a kind of representational content), the latter still fails to play any important explanatory role in the epistemic force of perceptual experiences.[60] I do so by considering a seldom remarked upon feature of blindsighted subjects.

iii. The Blindsight Argument for PA

Blindsight has been much discussed in both the philosophy of perception and the philosophy of mind more generally. This extraordinary condition occurs when all or a significant portion of the striate cortex is surgically removed or is otherwise severely damaged. It has been known for some time that the striate cortex is the primary seat for the processing of visual information by the brain. Any damage to or removal of this part

[60] There is a small caveat that must be addressed. Phenomenal contents, if they exist, can presumably be reflected upon and thought about. One can, if one has the relevant concepts, pick out a quale demonstratively. It is plausible, therefore, that one could come to warranted beliefs about the world in the following, somewhat irregular fashion: at t1 demonstratively identify a particular visual feel and associate it with state of affairs X, then at t2 identify the same (or a qualitatively indistinguishable) visual feel and infer that X holds. In such a sequence of reasons phenomenal contents could play an explicit role in coming to a perceptual belief about the world. I mean to ignore beliefs based on reasoning that makes explicit demonstrative reference to phenomenal contents, and am willing to grant that this is a way in which one could come to possess a warranted belief about the world. Such cases are different from cases of ordinary perception: it is clear that the warrant for the belief that X obtains is not *immediate*, because the belief is based on a belief that this very quale is associated with X. As such, there is a significant difference between such a case and the ordinary perceptual case in which one possesses immediate warrant.

of the brain is expected to cause blind areas (scotomas) in the subject's visual field corresponding to the removed or damaged portion of the cortex. In patients like DB and GY, two of the most carefully studied blind-sighted subjects, the absence of large portions of the striate cortex resulted in the predicted scotomas. Both patients deny seeing anything in a significant portion of half of their visual field, and they typically report no visual impressions or phenomenology from stimuli in the affected portion of their visual field. If you project an image in DB's blind field and ask him whether or not it is an X or an O, he will object that he can see nothing, that he has no idea what is there, etc. However, if you ask DB to *guess* whether the object is an X or O, DB will give the correct response at a rate far above chance, all the time maintaining that he is "simply guessing," that he has "no idea" what is there, and while denying that he has any visual experience at all in the relevant area of his visual field.

DB, GY, and other blind sighted subjects have demonstrated the ability to make many visual discriminations in this manner including (but not limited to) the orientations of objects, identifying the trajectory of a moving stimulus by pointing, distinguishing shapes (Weiskrantz 2009), and color discrimination (Stoerig and Cowey 1989). GY has even shown evidence of discriminating letters in his blind field and of words displayed in his blind field affecting his linguistic understanding of words later displayed in his normal unaffected visual field (Marcel 1998). A natural conclusion to draw from these surprising results is that, despite a complete lack of conscious awareness of any experience in the blind field, there is still some visual information being processed which results in a recoverable representation. The representation itself is not conscious, but

when a blindsighter is given a forced-choice task, the information informs their choice resulting in "guesses" that are highly reliable.

One might think that cases of blindsight support the claim that phenomenal contents play a crucial role in explaining our perceptual warrant. Indeed, Declan Smithies (2014, forthcoming), Siegal and Sillins (2011) and Silins (2010) all take blindsight to provide at least some support for views like PA*. Blindsighters after all, intuitively lack warrant for believing propositions about states of affairs in their blind field, and they quite notably lack visual phenomenology despite their reliable responses in forced-choice guessing tests. Of such subjects, one might reasonably conclude that their lack of perceptual warrant can be explained by the absence of any visual phenomenal contents. Here is a rough outline of how such an argument might go:

Blindsight Argument 1 (BS1)

1. A sighted subject has phenomenal contents for her visual experience while a blindsighted subject does not.
2. The blindsighted subject is not warranted (or not as warranted) while the sighted subject is.
3. They both reliably make true judgments about the objects in their field of vision.
4. The epistemic difference is best explained by the fact that the blindsighter lacks phenomenal contents while the ordinary subject does not.[61]

Obviously the argument as presented here is far from complete, but it is enough to get the flavor of the argument in favor of PA that arises from consideration of blindsight. First, we see that the argument assumes that blindsighters lack perceptual warrant, or enjoy less

[61] No supporter of PA of which I am aware offers this as an argument in favor of their view. I present it here in order to highlight the difficulties confronting any argument that attempts to argue from a blindsighter's lack of perceptually grounded warrant to PA.

warrant than similarly situated sighted subjects. This claim is not without plausibility, but it is far from unobjectionable. It is clear that since the blindsighter typically lacks beliefs about what is in her blind field, she lacks *doxastic warrant* in virtue of lacking beliefs. However, it is not so obvious that she is similarly devoid or deficient in *propositional warrant*. The propositional warrant may be unavailable to her in virtue of her because her deficiency restricts her access to the warranting perceptual representations. One would like some argument to the effect that is not the case. Nevertheless, I propose to grant this premise as sufficiently plausible to sustain an argument of some interest.

The more serious problem with BS1 is that it is unsound. One difference between blindsighters and ordinary subjects is the lack of phenomenology associated with the perceptual representation, but this is hardly the only difference. Here is a non-exhaustive list of other relevant differences:

1. Blindsighters are not aware that they are receiving visual information in the blind field.
2. They are not inclined to form beliefs about stimuli in their blind field (even when made aware of their reliability).
3. In many cases blind sighted subjects, though performing significantly above chance, are less reliable than control subjects, (or are less reliable at discriminating stimuli in their blind field than in their sighted field).
4. Blindsighters often take themselves to be guessing, rather than responding to actual information.
5. In some cases, they will claim that they must have done poorly when, in fact they were quite reliable, or claim to have done well when they have done poorly. Thus, they are incapable of discriminating good performances from bad ones in their blind field.
6. The information furnished by blindsight is not available for use in further reasoning or action.

All of these differences could play a role in explaining the epistemic differences between blindsighters and sighted subjects. Without considering the epistemic import of these further differences, we have no good reason to suppose it is a lack of phenomenology that explains blindsighters' lack of perceptual warrant; at least some of the other differences listed above have at least *prima facie* plausibility as culprits.

Smithies (2014) attempts to fill in this gap in the above argument by constructing a series of thought experiments based on actual cases of blindsight that attribute progressively more robust cognitive capabilities to their subjects. At each stage he argues that, even with these additions, it is still implausible that the subject enjoys ordinary perceptual warrant for her beliefs. Eventually, he takes himself to have eliminated any epistemically relevant cognitive difference between a blindsighter and a normal perceiver excepting only the lack of phenomenology. He concludes that the epistemic difference between blind-sighted subjects (actual and imagined) is best explained by the blindsighters' lack of phenomenal character.[62]

Smithies begins his discussion by asserting the efficacy of two counterexamples to a very basic version of reliabilism. This version of the theory has it that reliability (suitably cashed out) is both necessary and sufficient for warranting belief. Smithies claims that both theses can be shown to be false by well-known thought experiments.

[62] Smithies proposes a general account of epistemic warrant. That is, he extends his claims to *all* cases of warranted belief, not just cases of basic beliefs, or perceptual beliefs. I will reformulate various theses and claims that he makes to make more limited claims about perception and the typical way it warrants basic beliefs. If his claims about perceptual warrant are false, then his general claims will also be shown to be false, however it will not follow that the view he advocates is not acceptable for any kind of warrant. That is, even though I contend that his view is an inadequate account of perceptual warrant, I remain (as always) neutral as to whether his claims could provide an adequate account of other varieties of warrant (e.g. introspective, *a priori*, memorial, etc.)

Envatment

My envatted mental duplicate has warrant to form beliefs on the basis of perceptual experience, memory, testimony, and so on, although forming beliefs in this way is unreliable in the circumstances. (Smithies 2014)

Clairvoyance

My clairvoyant mental duplicate lacks warrant to believe on the basis of blind hunches, wishful thinking, and so on, although forming beliefs in this way is reliable in the circumstances. (Ibid.)

The first case has its canonical presentation in Cohen (1984) and the latter is due to BonJour (1985). Envatment is intended to establish that reliability is not necessary for warrant; my envatted duplicate is warranted while forming beliefs on the basis of unreliable belief-forming processes. Clairvoyance is intended to establish that reliability is not sufficient for justification; a BonJour-style clairvoyant is maximally reliable but not warranted in her beliefs.

A few comments are in order: first, while I share the intuition appealed to in Envatment, I think that the case is far more complex than Smithies' appreciates. In the last chapter of this dissertation, I will argue that the case surreptitiously relies on assumed cognitive and historical similarities between myself and my vat duplicate that make it plausible that even while being radically unreliable he is, nonetheless, warranted. In considering cases where my "duplicate" lacks these similarities, the intuitive force of the case is far less clear. I refer the reader to Section 8 of Chapter IV for a more complete discussion of this issue.

Second, the conclusion that Smithies' infers from Clairvoyance would not and should not be granted by any philosopher who considers herself to be a reliabilist. BonJour (1985) presents a sequence of clairvoyance cases, the clearest of which builds in that the clairvoyant subject has significant contrary and/or defeating evidence that

intuitively undermine her clairvoyant beliefs. A very implausible kind of reliabilism that cannot take into account the presence of defeating evidence will be committed to claiming that, since the clairvoyant's beliefs resulted from a reliable belief-forming process, he is warranted even while entirely ignoring a host of defeating evidence. However, in the case in which the subject has no defeating or overriding evidence, and her clairvoyant faculty and reliability is understood in its intended sense (not simply "getting it right" frequently, but as resulting from a counterfactual supporting capacity) her miraculous ability begins to look more and more like any other basic belief forming method (perception, introspection, etc.). The reliabilist could simply accept as unproblematic the conclusion that she is warranted in forming beliefs on its basis.[63]

It is also important to note what Envatment and Clairvoyance do not show, even if one accepts their intended conclusions. Showing that something is neither necessary nor sufficient for epistemic warrant is a long way from showing that it is epistemically irrelevant. If reliability is not necessary for warrant generally, it could still be an important consideration for some varieties of warrant. Even if one could show that reliability is not necessary for *any* particular sort of warrant (e.g. perceptual warrant), it could still be the case that reliability could be a part of some jointly sufficient conditions for warranted belief. More importantly, Envatment at best shows that a kind of *local* reliability is not necessary for perceptual warrant. The BIV is not reliable in the environment in which she finds herself, and if she is still intuitively warranted she is so even though she is massively perceptually deceived about her environment. However, if

[63] Thanks to Tony Breuckner for many helpful discussions on these issues. I return to some of these points in Chapter IV.

one maintains that there is a privileged environment for the evaluation of the reliable veridicality of one's perceptual representations, then the fact that a subject finds themselves in an abnormal environment in which her perceptual systems are unreliable does not impugn her warrant.[64]

The purported conclusion of Clairvoyance, that reliability is not sufficient for warrant, is particularly weak. One can consistently think that forming a belief using a reliable belief-forming method is not sufficient *on its own* to produce warranted belief while thinking that it is one of several conditions that, when all are met, result in a warranted belief. On such views, the reliability of one's perceptual system is an important element in one's perceptual warrant but it is not the *only* relevant consideration. Indeed, I take it that there are very few who would advocate the view that reliability veridicality of one's perceptual representations is *sufficient* for warrant, but many who have convinced that it plays some necessary role in warranting perceptual beliefs. BonJour-style cases do not present significant problems for such views.

However, let us put these complaints aside and consider Smithies' argument for PA*. He argues that it is the lack of phenomenal content that explains the blindsighter's lack of perceptual warrant by considering two counterfactual blindsight cases in which the subjects have more and more at their disposal, cognitively speaking, but still lack conscious phenomenal experiences. Such subjects, Smithies, claims, lack warrant for their beliefs. Since adding additional cognitive resources to blindsight cases while leaving out phenomenal consciousness does not result in warranted belief, he concludes

[64] I will defend an account with this feature in Chapter IV.

that what accounts for the epistemic failings of blindsighters is their lack of phenomenal content.

Blindsighters like DB and GY are not at all disposed to form beliefs about objects or phenomena in their blind fields. If asked about what is in their blind field, they will typically claim to see nothing at all. Even when the subjects can be persuaded to make a guess about what is happening in their affected field[65] they claim to just be choosing at random and are (initially) surprised when informed of their success. They claim to be able to see nothing in that area of their visual field, and as such are not disposed to form beliefs about what is there. Even after being informed of their reliability, the information that allows them to succeed in forced-choice guessing tasks is not available to them in their normal cognitive activities. They are not inclined to form beliefs about objects or occurrences in their blind field and, hence, cannot make use of the information in carrying out actions, plans, or in fulfilling their desires. They may know from engaging in the experiments that they have surprising discriminative capabilities in their blind field, but those capabilities only manifest themselves in ways that make the information doxastically and practically useless. In most every day circumstances, they are operationally blind in the affected part of their visual field.

Given these differences between the blindsighter and an ordinary perceiver, the following explanation for their epistemic differences is plausible: the reason that the blindsighter lacks warrant is that the visual information is not *accessible* in the same way

[65] Some subjects simply refuse to offer a "guess." After all, they are being asked to identify an object in an area of their visual field in which they have no conscious visual experience at all. Imagine if someone asked you to guess at whether a piece of paper, held just behind your head held an "X" or an "O." One can imagine being exasperated and refusing to participate in so ludicrous a study. Indeed, one subject tested by Weiskrantz categorically refused to participate exclaiming "I do not gamble sir!"

as it is to a sighted subject. The blindsighter's reliability in forced-choice guessing tasks seems to show that visual information is being processed *somehow*, but it is clearly not accessible in a way that would allow for one to form beliefs on the basis of that information. One might think that if the information were to be made accessible to the blindsighter, then perceptual warrant would follow. Such availability need not entail any conscious phenomenal experience, rather, one might think, all that would be needed would be for the blindsighter to be capable and inclined to form beliefs on the basis of the residual visual information that informs his forced-choice guessing success.

Smithies (2014) argues that this contention is shown to be false by considering Block's case of the super-blindsighter. A super-blindsighted subject is just like a blindsighted subject except that she is "disposed to use unconscious perceptual information in the direct and spontaneous control of thought and action without any need for prompting" (Smithies 2014). In other words, the super-blindsighter is inclined to form beliefs, plan and carry out actions directed at, and make assertions (without prompting) about, the objects in her blind field. It is useful to consider a concrete example: the super-blindsighter confidently forms beliefs that, for instance, there is a pencil in front of her, when the pencil is presented in her blind field, she will reach for it if she needs a writing implement, and if asked where one could find a pencil will confidently point or otherwise indicate the location of pencil. If one were to replace the pencil with a cup of water, she would reach for the cup of water if thirsty, inform someone whom she knew to be thirsty of its location (or move to hand it to them), correctly and confidently assert that there is a cup of water there and so on. Still, in this case, she insists that she is not conscious of the objects in her affected field. She may be

mystified as to how or why she is so confident about her belief and somewhat surprised that her confidence is reliably confirmed. There is nothing that it is like for her to be experiencing the pencil and cup, but she nonetheless forms beliefs and is disposed to act in the manner described.

Despite these remarkable capacities, Smithies states that "the super-blindsighter is in the same epistemic predicament as the blindsighter," because, "the mere feeling of confidence is not sufficient to justify forming beliefs—justification is not that easy to come by!" (Smithies 2014). Smithies' reasoning here is unsound. One could grant that mere confidence is not enough to guarantee a warranted belief, but that is not all the super-blindsighter's belief has going for it. It is a true belief, formed in a reliable fashion, by use of a perceptual faculty. The subject is disposed to make use of the information in successful interactions with the objects in her blind field, and to report correctly as to their presence or absence. At this stage of the argument I do not insist that the super-blindsighter has a warranted perceptual belief,[66] I simply wish to suggest that it is far from clear that she is *unwarranted*. Smithies seems to assume that if confident belief formation on its own is not sufficient for the production of a warranted belief, then adding it to a subject's capacities in a hypothetical case will never produce a case in which the subject enjoys a warranted belief. But this is fallacious. That a certain property X is not sufficient on its own for satisfying some condition C, does not entail that X is not part of some jointly sufficient set of properties that together guarantee that C

[66] I find the claim that the super-blindsighter is warranted in her perceptual beliefs very plausible, but at this stage I simply wish to point out that nothing Smithies says supports the conclusion that she *isn't* and thus he fails to marshal support for PA by considering the case of blindsight.

is satisfied. This argument from Smithies fails to establish the conclusion that the super-blindsighter is not warranted.

Smithies attempts to support his judgments about this case by comparing the super-blindsighter to BonJour's clairvoyant:

> In effect, the only relevant difference between blindsight and super-blindsight is the addition of a reliable doxastic disposition, but as the clairvoyance case illustrates, the mere fact that beliefs are formed in a reliable way is not sufficient to make them justified (Smithies 2014).

> The epistemic predicament of Susan the super-blindsighter is closely parallel to that of BonJour's clairvoyant Norman…Intuitively Norman has no justification to form beliefs about he location of the President on the basis of his reliably faculty of clairvoyance. By parity of reasoning, Susan has no justification to form beliefs about objects in her blind field on the basis of her reliable faculty of unconscious perception (Smithies forthcoming).

I have already suggested that Smithies should not be nearly so sanguine about his conclusions regarding Clairvoyance as he is; but let us simply grant his diagnosis for the sake of argument. Even with this substantial concession, Smithies' reasoning here is still problematic for the same reason indicated above; the clairvoyance cases show, at most, that reliable belief formation is not sufficient for warrant on its own. They do not establish that reliable belief formation is irrelevant to a subject's warrant. Reliability could, for all BonJour's cases show, still be a necessary condition for certain kinds of warrant. In particular, having a perceptual system that is reliably veridical could be one of several conditions that must be met for a subject to form warranted immediate perceptual beliefs. Showing that reliable belief formation is not sufficient for perceptual warrant does not establish that reliability isn't one of a set of jointly sufficient of conditions for warranted belief.

Even if we ignore these problems and grant Smithies' conclusion that the super-blindsighter lacks warrant for her perceptual beliefs, this still does not give us persuasive reasons to accept PA*. As Smithies correctly notes, there are still differences between an ordinary perceiver and a blindsighter, other than the latter's lack of phenomenology, that could be used to explain the epistemic differences between them.. In particular neither the blindsighter nor the super-blindsighter form higher-order thoughts and beliefs about the perceptual information in a "non-inferential and non-observational way" (Smithies 2014). Smithies seems to have in mind beliefs such as: "it perceptually appears to me as though p," "I am having a visual experience with the content of p," and so on. As he correctly points out such capacities are not necessary for perceptual warrant because small children and higher animals have warranted perceptual beliefs but lack the kind of introspective capacities and/or cognitive complexity that would allow them to form beliefs about their own perceptual experiences.

It is notable, however, that Smithies maintains that even if one were to *grant* the blindsighter capacities of this kind, she would still lack warrant for her perceptual beliefs. He argues for this conclusion by considering the case of *hyper-blindsight*: "which is just like super-blindsight except that the subject has a reliable disposition to form higher-order thoughts about unconscious perceptual information in a non-inferential and non-observational way" (Ibid.). The hyper-blindsighter, according to Smithies, is no better off than the blindsighter: all that has been added is a reliable disposition to form higher-level beliefs, but since mere reliability does not suffice for warrant, these higher-level beliefs will themselves be unwarranted, and one does not gain warranted perceptual beliefs merely by adding more unwarranted beliefs to one's doxastic profile.

To assess the plausibility of Smithies' contentions about the hyper-blindsighted subject, let us consider the abilities of such a subject in more detail. Smithies does not clearly define what is requires to be "meta-cognitively conscious" of the unconscious perceptual information on the basis of which her belief is formed. Suppose, however, that the hyper-blindsighter is disposed to reliably and non-inferentially form beliefs that e.g. she belief that there is a red cube in front of her because she *sees* it. However, since she is blindsighted, she will deny having any phenomenology or conscious visual experiences of the things in her blind field. Let us imagine what happens when we place a sphere in the blind field of such a subject. She will immediately and confidently form the belief that there is a sphere there, as does the super-blindsighter. Furthermore, she will also immediately and confidently form beliefs such as: "I am having a visual experience of a sphere," "I got my belief that there is a sphere here because of my (unconscious) visual experience as of a sphere," "it appears visually to me as though there is a sphere here" and so on. The hyper-blindsighter not only correctly and reliably forms beliefs about the objects in her blind field, she *also* correctly and reliably believes that her beliefs are the result of her (unconscious) visual experiences.

Smithies maintains that these higher-level cognitive capacities cannot contribute to the warrant for the hyper-blindsighter's perceptual beliefs because, lacking phenomenal basis, they themselves cannot be warranted; one does not, he claims, arrive at warranted beliefs simply by adding further unwarranted beliefs (Smithies 2014, forthcoming). With this I agree. However, Smithies argument that the hyper-blindsighter lacks warrant is inadequate. First, it is far from obvious that the hyper-blindsighter's meta-cognitive beliefs are unwarranted. My position of strict neutrality on

non-perceptual epistemological issues in the context of this dissertation entails that I have no theory of introspective warrant to offer. However, I am loath to simply grant that some variety of phenomenology is necessary in order for one to be warranted in believing that one is having a particular sort of mental event like a perceptual experience. Smithies' contention seems to be grounded in his claim that phenomenology is necessary for warrant *generally*, but one can hardly rely on such a claim in attempting to establish that very conclusion.

Secondly, from a strictly intuitive standpoint the claim that the hyper-blindsighter is unwarranted in her perceptual beliefs seems quite remarkable. If I reflect on how I would judge the epistemic status of the hyper-blindsighter were I to meet her, I cannot help but think that I would attribute warranted perceptual beliefs and, indeed, perceptual knowledge to her. I would do so based on her similarities to other subjects whom I take to be warranted in their perceptual beliefs: she confidently and accurately forms beliefs about objects in her blind field, and can use those beliefs in carrying out actions and in furthering her plans. Furthermore, if asked *how* she knows that, e.g. there is a ball in front of her, she will respond much as I and other normal adult perceivers would, that she *sees* it and/or has a visual experience of the ball. If I were to subsequently find out about her peculiar condition, I do not think it likely that I would retract my former evaluation of her epistemic status. I would conclude that, despite her condition she warrantedly believes that there is a ball in front of her.

I do not make these observations in order to argue directly against PA*. It is plausible that the conditions for warrant *attribution* do not provide an infallible guide to the necessary and sufficient conditions for warrant *possession*. What these reflections do

suggest is that the claim that the hyper-blindsighter lacks warrant for her beliefs is of questionable plausibility. One would like a persuasive argument to the effect that this is so, but, as yet, none has been offered. It seems, then, that Smithies' argument, and blindsight cases more generally do little to support the phenomenal approach. In the next section, I argue that quite contrary to Smithies' conclusions, blindsight cases actually present a serious problem for PA*.

iv. Uber-Blindsight and the Phenomenal Approach

Kentridge and Heywood point out that the philosophical interest in blindsight patients stems mostly from their status as "real-world zombies" (Kentridge and Heywood 1999, 4); that is, a subject who has representational states without any corresponding phenomenology. Indeed, the striking ability of such subjects to make highly reliable discriminations while adamantly denying any awareness or conscious experience is of enormous interest for both epistemology and the metaphysics of mind. However, the surprising and counterintuitive nature of this feature of the condition has caused many to overlook other important elements in the reports of such patients. It is seldom remarked upon that at least some blindsight patients *do* sometimes report having experiences, or of being aware of stimuli in their blind fields. What is of particular interest to our discussion of PA* is that these experiences can be of a highly aberrant nature.

Wieskrantz distinguishes two varieties of blindsight:

> **Type-1 Blindsight:** The ability of a subject to make reliable visual discriminations without any conscious awareness reported of the stimuli.
>
> **Type-2 Blindsight:** The ability of a subject to make reliable visual discriminations while reporting awareness of something presented but claiming not to "see it." (Weiskrantz 2009, 14)

The blindsighted subjects that I know of all experience instances of both Type-1 and Type-2 blindsight. In some cases of Type-2 blindsight the reported awareness amounts to little more than a certainty that "something" is happening in the blind field,[67] but this is not the only kind of awareness that blindsighted subjects report. In some cases they report, not just awareness of some stimuli, but robust and radically aberrant phenomenal experiences. In most cases DB reports no awareness of any experiences or stimuli and is significantly surprised when he is informed of the reliability of his "guessing." For instance, when, after one trial DB was asked whether he had known he was correct nearly every time he replied: "No, I didn't—because I couldn't see a darn thing" (Weiskrantz 2009, 87). However, in some cases, especially when the stimuli were either of very high luminance, long duration, or rapid onset or movement, DB would give noteworthy reports of his experiences.

DB made the following reports in a test in which he was asked to point, with an outstretched finger, to stimuli in his blind field (he succeeded at a rate far above chance). DB denied having visual experiences of any kind, but made the following reports:

> "They appeared to stand out in front of the screen. I felt I could push them back."
>
> "The stimulus stood out from the perimeter about 3 inches."
>
> "I have a feeling of a quick movement out to about 45 degrees."
>
> "There was definite movement. Something seems to pop out a couple inches...I didn't see any light. If anything it was sort of dark. A moving wave."

[67] In some such cases (but not all), the awareness can be plausibly explained by the subject's (perhaps learned or acquired) awareness of saccadic eye movement or perhaps even of his pupil dilation (Cowey 2004). In such cases the awareness would not be visual, but rather some learned discriminatory awareness of muscle contractions in the cuing the subject to the presence of a stimulus. Such cases are of little interest here.

> "I felt as if something was coming up to me. But I didn't *see* anything." (Ibid. 104-105)

During a test involving a moving vertical line DB described his experience as follows: "I did not see the line. There was a kind of pulsation. I did not see the light or anything, but I could *feel* the movement and I was absolutely sure of it" (Ibid. 127). In a test attempting to establish whether or not DB could discriminate the orientation of an object in his blind field, he seems to instinctively correlating his peculiar experiences with particular orientations of the objects:

> "I *think* I say "across" when the bolt coming out is thinner, and "not across" when it is thicker. But maybe not. It is somewhat like what my moving finger looks like when I close my eyes in my imagination" (Ibid. 133).

> "One is definitely thicker than the other--I say "across" when it is thin, but I don't know whether I am right or not" (Ibid. 134).

Similarly in a test in which he was asked to guess whether or not the object in his blind field was an X or O, DB would deny seeing anything, but say he could "feel" something approaching or receding from him. When the approaching or receding object "felt smooth*"* he said "O," when it "felt jagged" he said "X." (Ibid. 93) In some cases of vigorously moving stimuli, DB claimed that he had a genuine visual experience, but his experience would be of a radically aberrant variety: rather than a small moving stimuli, DB reported "complex patterns of radiating lines and grids." (Ibid. 123)

Most striking, however, are DB's experiences of what he calls 'waves' in response to vigorously moving or highly illuminated stimuli in his blind field. Here is Weiskrantz's description of this phenomenon:

> The experience is of a kind unlike anything in normal visual experience for which precise words seem to be lacking. The 'waves' can have some sort of form. They can be straight or curved, or can even have 'squareness.' They can also be 'quick.' The waves could also be thinner or thicker, e.g. in the uncurved vs. curved triangle discrimination…(DB said, describing his experiences) "there is movement. I said 'straight' when the waves are thinner, and 'curved' when they are thicker and quicker." But D.B. insisted that he was not '*seeing*.' (Ibid. 160)

Even in cases in which DB has aberrant phenomenology, the other characteristic features of his blindsight are still present: he typically denies seeing anything, he is not inclined to form beliefs on the basis of his experiences, he maintains that he is just guessing about what is in his blind field, and so on. As one might expect, having an experience does sometimes cause DB to be more confident that he has performed well in a particular trial, but this is not always the case, but it is not always the case that his having such an experience correlates with his good performance. In particular, in the case of the "waves" described above would often seem to *hamper* DB's discriminative capabilities in his blind field. This could happen when the differences in the experienced waves did not covary in the proper way with the stimuli of the test, or in which they constituted a distraction causing him to attend to the waves rather than to the stimulus (Ibid. 161). Even in cases in which DB was highly reliable in his discriminations while having an experience, Weiskrantz and colleagues were often able to show that DB was capable of making reliable discriminations even in its absence,[68] showing that his discriminatory capabilities do not depend solely on his experiences.

[68] They usually accomplished this by reducing the luminance, the suddenness of the onset of the stimulus, or the rapidity of its movement. Having once confirmed that DB now had no conscious awareness of the stimuli, they ran the test again, with the result that he was still reliable. In most cases, he remained as reliable, or was even more reliable than he was when he was conscious of some experience.

GY, another subject with massive damage to his striate cortex often exhibits blindsight responses of the Type-1 variety.[69] However, in some cases he also reports conscious experiences of stimuli in his blind field. Unlike DB, however, GY's conscious experiences seldom hamper his ability to reliably discriminate amongst stimuli in his affected field. Indeed, GY's awareness of stimuli in such cases caused Barbur et. al. (1993) to conclude that, GY was not blindsighted in the standard usage of the term:

> Whether tested subjectively—through a verbal report or objectively—through discrimination, the subject gave every sign of having seen and having been consciously aware of what he had seen…GY verbally reported seeing movement…he gave a verbal report of the direction of motion. (Barbur et. al 1993, 1295)

However, in an interview conducted by Weiskrantz about the nature of his experiences during this very same study, it is revealed that GY's experiences are of a curious and largely indescribable character:

> LW: What kind of words do you think you have to use (to describe the experience)?
>
> GY: The nearest I ever get, and it is not a fair comparison, is waving your hand in front of your eyes while they are closed. You are kind of aware that something has happened, but you don't quite see it. You know something has moved. But that isn't a fair comparison to make.
>
> LW: The nearest you can get is the sense of something happening, but you don't know what it is?
>
> GY: Yes.
>
> LW: Anything else you can think of to describe it?

[69] Weiskrantz, Barbur, and Sahraie (1995) tested GY extensively to determine the conditions under which GY had conscious experiences and whether, in the absence of such experiences, his discriminative capacities remained. They found that GY was most likely to report experiences in cases of high luminance contrast, and high movement speed of the stimulus. GY was still able to perform well above chance even when conscious experiences were absent.

> GY: No, because it is a sense that I haven't got, if that makes sense. If you said something to try to describe sight to a blind man, we don't have the words to do it because he does not have the receptors or the reception, and that is the same with me. I mean I can't describe something if I don't understand it myself. (Weizkrantz 1997, 144-145)[70]

Given these various reports of blindsighted subjects of experiences beyond simple awareness that something or other is happening, it is useful for our purposes to distinguish a third variety of blindsight:

> **Type-3 Blindsight:** The ability of a subject to make reliably veridical visual discriminations while reporting both awareness and phenomenology, but while reporting the phenomenology as non-visual or otherwise radically aberrant.

It is far from clear how we ought to epistemically evaluate DB and GY in cases in which they have some sort of conscious awareness of stimuli in their affected field. In DB's case, I am inclined to the view that he lacks even propositional warrant for believing anything on the basis of his aberrant experiences. Even though having such experiences can increase his confidence in particular judgments, they don't seem to covary with the target stimulus, and, in many cases, interfere with his discriminative capacities.

The case of GY is less clear. While his experiences are clearly of a peculiar nature, they do allow him to confidently and reliably make judgments about stimuli in his affected field. However, in the reports cited above, it is unclear whether or not GY's is genuinely a case of Type-3 blindsight. He reports awareness of *something* happening, and finds this awareness difficult to describe, but it is not clear that he has any of the sort of robust qualia described by DB in the cases above. GY's inability to give words to

[70] Barbur et. al. concede later in the same article that GY's expriences are "not the same as that of normal individuals" (Barbur et. al. 1993, 1301). Thus, their report about the test in question do not conflict with GY's later reports.

experiences makes their character difficult to ascertain. There is some empirical evidence that he will at least sometimes have experiences with phenomenal character above and beyond a simple awareness of something happening, but these results are inconclusive.[71] It could be that his reluctance to claim that he "sees" in many such cases is a result of radically different but still robust phenomenology of the kind reported by DB. On the other hand, it could indicate that his experience is merely an awareness of *something* happening and thus an instance of mere Type-2 blindsight. Like DB, GY is typically not inclined to form beliefs even when he has conscious experiences of stimuli in his affected field, nor to utilize the information that affects his forced-choice guesses in his everyday actions, reasoning or decision-making. For these reasons, if no other, I am inclined to conclude that GY like DB lacks even propositional warrant to believe anything on the basis of his perceptual experiences. However, it would be a dubious procedure at best to draw firm epistemic conclusions about such cases, and I will not attempt to do so here. I suggest, therefore, that we use these cases as a conceptual crutch, abstract away from the personal peculiarities of DB and GY, and construct cases that we can, perhaps, judge on a firmer intuitive basis.

Let us consider another hypothetical subject: call her an *uber-blindsighter*. The uber-blindsighter is like Smithies' hyper-blindsighter, but her blindsight is always of the Type-3 variety. That is, she always has some sort of phenomenology associated with

[71] Stoerig and Barth (2001) had some success in getting GY to report something close to a phenomenal match between his experience in his blind field, and a projected a distorted image of the original stimulus in his unaffected field, which suggests that his experience is more than simply an awareness of something happening. However the results were far from conclusive. They were not successful in matching many of his experiences because in the experience in his unaffected field was described as "merely visual" and thus "no match at all." (581)

retinal stimulations in her affected field. However, unlike actual blindsighted subjects, she will reliably and confidently form immediate beliefs about objects in her affected field. To make the case as strong as possible let us also stipulate that she will reliably, confidently, and immediately form beliefs that she is having visual experiences and what those experiences represent, and her experience is always accompanied with a phenomenology, though of a radically aberrant nature.[72] She is, like DB, aware of the difference in phenomenology between her affected field and her ordinary field. She will report, for instance, that when there is a square placed in her affected field, her conscious experience is one of shifting waves of some kind or other, difficult to describe, while in her normal field she reports the phenomenology we would expect of a normal subject. Nonetheless, she can make comparative judgments across the fields (e.g. reporting that the shape in her affected field is the same as that in her normal field, she will track an object moving from one field to the other, notice changes and so on).[73] Assume that all such judgments are not based on any background beliefs that the uber-blindsighter possesses and highly reliable. Finally, when we ask her *how* she makes judgments about objects in her blind field she will, as an ordinary subject would, be inclined to point to her

[72] I anticipate the objection that the uber-blindsighter is not really a blindsighter at all. A blindsighter, one might think, must at least in most cases fail to have any conscious experience of objects in her blind field. I am happy to concede the taxonomical point. However, this will be irrelevant to my objection to PA. All that will matter is that the case that I describe is possible, the word we use to describe the condition of the subject is immaterial.

[73] Interestingly, there is some evidence that actual blindsighters are capable of reliable comparative judgments of just this kind. In several tests, DB performed well above chance when tasked with determining whether or not the stimuli in his affected field was the same as or different as that in his unaffected field. There is also some evidence that stimuli in the affected field of a blindsighted patient can affect judgments of stimuli in the subject's *unaffected* field (Marcel 1998).

experiences. For instance, she may make reports of the following form (again analogous to DB's): "when the waves are such and such a way, I say and believe 'square' or 'red.'" From her perspective, such experiences are distinguishable and each associated with a particular representational content.

I find it quite plausible that the uber-blindsighter is warranted in her perceptual beliefs about objects in her blind field, but this is not my primary contention. The main thing to note about the uber-blindsighter is that an *internalist,* in the sense I have defined the term, should have grave reservations about denying that the uber-blindsighter possesses warrant for her perceptual beliefs. From her perspective it is no accident that she has formed the relevant belief, she forms it on the basis of a conscious visual experience that she believes that she is having and believes represents the object as being the way she believes it to be. Moreover, she is stipulated to form these beliefs in a non-inferential manner in much the same way as she does about objects in her unaffected field of vision. Of course, as a reflective agent she can think about her experiences, report on the differences, and so on, but typically such higher-order capacities do not play a role in the formation of her perceptual beliefs whether in her normal or affected field. However, it is clear that the basis for her beliefs, the aberrant conscious experience, is introspectively available to her. She is aware of the difference between her affected and unaffected fields, but does not regard this difference as any defeater on her beliefs. The aberrant experiences cause her to form confident perceptual beliefs and she is not at all inclined to doubt that they are veridical.[74] From an internalist perspective there seem to

[74] To make this point more salient, let us suppose that her condition was not the result of some accident or medical procedure, but is the result of a congenital defect. For her things have *always* been this way.

be no solid grounds for denying that the uber-blindsighter is warranted in her visual beliefs.[75]

Furthermore, unlike the Inverted Spectrum Case, the uber-blindsighter is compatible with the weak version of representationalism defined in sub-section ii. There is a key difference between the uber-blindsighter and the Inverted Spectrum. In Inverted Spectrum, Andy and Ollie were capable of having experiences with the same phenomenal character that warranted different beliefs, but the uber-blindsighter's experiences do not have this feature. Rather, her experiences in one part of her visual field are *radically phenomenally aberrant,* in comparison to those in the other. Indeed, we may assume that there is no phenomenal similarity whatsoever between the areas of her visual field. In other words, it is entirely consistent to maintain that the contents of uber-blindsighter's visual experiences are phenomenal contents as defined above, and hence the case is, strictly speaking, compatible with both representationalism and PA*.

In the uber-blindsighter we have a subject who has radically different phenomenal experiences in different parts of her visual field. As I have described the case, the experiences do not share any intrinsic phenomenal similarities, and yet they succeed in warranting the same perceptual beliefs. Recall that PA* is not just the thesis that one must be *conscious* in order to have warranted perceptual beliefs. The uber-blindsighter

[75] An externalist who thinks that *proper functioning* of one's perceptual system is necessary for perceptual warrant may have cause to demur from this assessment; there is clearly something abnormal about the subjects visual processing and that may support the judgment that the experiences are the result of a significant sort of improper functioning in her visual system. However, even here the case is not entirely clear. One interpretation of what occurs in actual blindsight patients is that such subjects' discriminatory capacities are explained by visual processing that typically happens outside of the striate cortex. According to this hypothesis, then, the mechanisms responsible for the visual processing that remains may, in fact, be functioning properly.

meets this condition. Rather, the proposal was that something about the *particular* phenomenology of perceptual experiences explains their warranting force. But if a single subject can be warranted in believing the same propositions based on viewing the same visual scene in the same conditions, based on radically different phenomenologies, one requires some explanation of this fact. Call the particular phenomenal character in the uber-blindsighter's affected field corresponding to the content that there is a square here under such and such lighting conditions, at a certain distance, etc., A*. Call the corresponding phenomenal character of her experience in her normal field with the same content N*. The proponent of PA owes us some explanation for how both A* and N* are capable of warranting beliefs with the same contents in the same subject. Given that A* and N* share no intrinsic phenomenal similarities, it is unclear how this may be done without appealing to non-phenomenal properties of the experience.[76]

I anticipate two replies on behalf of PA*. As I have already conceded, a phenomenal content is a content that is shared by any two experiences with the same phenomenal character. This is compatible with two experiences with different phenomenal *characters* sharing the same phenomenal *contents*. Therefore, a defender of PA* can maintain that the warranting force of A* and N* *is* explained by phenomenal contents that both experiences share. In other words, PA* by appealing to phenomenal contents is only committed to claiming that *if* two experiences share the same phenomenal character, *then* they will warrant the same beliefs.

[76] One possibility here is to maintain that there *is* an intrinsic phenomenal similarity between the two experiences: namely that both experiences have the *feel* of truth. In other words, there is a phenomenally individuated *seeming to be true* in virtue of which the experiences, though otherwise phenomenally different, can both warrant beliefs. This view collapses to a kind of phenomenal conservatism, and, as such, I put off discussion of this issue until the next section.

One could also seek to defend PA* by insisting that in both the case of the uber-blindsighter and the sighted subject, their phenomenology explains their warrant. Of course, the explanation for each subject's warrant is different, in the sense that each subject's warrant is explained by a different phenomenal experience. Nonetheless, one may insist, particular phenomenal features of each subject's experience suffice to explain the warrant they have for their perceptual beliefs. So the uber-blindsighter's warranted belief that there is a square there is explained by A* while the sighted subject's warranted belief is explained by N*. In each individual case the warrant is explained solely in virtue of the phenomenal content of the subject's experiences, though the phenomenal features of each are different. PA* could be defended by admitting that the *same* phenomenology is not necessary for different subjects to be warranted in the same beliefs, but that in each particular case the warrant is explained by phenomenal features of the experiences alone.[77]

These responses do little to blunt the force of my objection. Both amount to the substantial concession that the phenomenal features A* and N* have little to do with explaining why they warrant their respective perceptual beliefs. This seems to me to be a hollow victory. According to the first proposal on behalf of PA*, an experience with a certain phenomenal character guarantees that the experience has particular content and can warrant beliefs with that same content while remaining neutral about the cases in which the phenomenology of a perceptual experience differs, but the content remains the same. However, if two such experiences both warrant the same belief, then the phenomenology itself seems irrelevant to explaining the sameness of warrant. One can

[77] Thanks to Michael Rescorla for articulating this response on behalf of the phenomenal approach.

insist that they do share the same *phenomenal contents* and that is all that PA* requires, but if this is so it seems that it is not in virtue of any *phenomenal* similarity. The explanatory load must be carried by other factors, for instance: sameness of causal history, functional similarities, sameness of content-fixing relations, constitutive connections between particular phenomenal characters, and their contents, and so on. But if these properties of perceptual experiences must be appealed to in order to explain their warranting force, it seems as if we have abandoned the central motivation of the phenomenal approach. We ought to carry out our explanation in these terms alone rather than insisting on keeping the phenomenal as a spandrel in our explanatory architecture.[78]

A similar problem afflicts the second response. One can insist that there is a different explanation for the uber-blindsighter's and the sighted subject's warrant only by incurring a significant explanatory burden: namely, to say what it is about A* and N* that allow them to warrant the same beliefs. Since these two experiences, by stipulation, have no phenomenal similarities, it seems that this explanation must appeal to non-phenomenal features of the experiences. Thus, while the proponent of PA could consistently maintain that the phenomenal contents of A* and N* explain the warrant their respective subjects

[78] Smithies (forthcoming) amends his view to allow that non-phenomenal features of a perceptual experience are epistemically relevant only insofar as they pertain to the fixation of the contents of these experiences. Thus, his new view, unlike the old, is consistent with content externalism. However, the point above makes clear that this new position is still subject to the objection I have presented here. The uber-blindsighter and the sighted subject are stipulated to have the same perceptual contents. Smithies' new position maintains that once one has fixed the contents, the phenomenology suffices to explain a subject's perceptual warrant. But once again, this provides us with no answer to the question of what the *particular* phenomenal features of A* and N* contribute epistemically. Since they share no phenomenal similarity, the contribution must arise from non-phenomenal features. If there is some very general feature of such phenomenal states (e.g. a "truthy feel") then Smithies' view collapses into a kind of phenomenal conservatism. I consider the prospects for such a view in the Section 4.

enjoy for their perceptual beliefs, it provides no explanation for why these two entirely distinct phenomenal states can suffice to warrant the same beliefs. Once again, we would be forced to consider non-phenomenal characteristics of the experiences in order to explain their similar epistemic roles in their respective subjects. In either case, it seems that PA* has, at best, shifted the challenge posed by the meta-epistemological question back a step. To address this renewed challenge, we must appeal to non-phenomenal features of the perceptual experience.

v. Reductionist Representationalism

The arguments of the last section may suggest to some that the problem is that the connection between phenomenology and representational contents I have been making use of in my discussion of PA* is too weak. Instead of endorsing the merely conditional claim that if two experiences have the same phenomenal character, they must have the same contents, the proponent of the phenomenal approach ought to adopt a more radical representationalist view of the relation between phenomenology and contents. According to such a view phenomenal characters are identical or reducible to representational contents.[79] In other words one could endorse the proposal that one's experience has a certain phenomenal character *just in case* it has a particular (very complex) representational content. According to such a view, the uber-blindsighter's experiences with phenomenal characters A* and N* cannot share the same contents (at least not

[79] Smithies (2014) takes just this position. He rejects this position in his more recent work (Smithies forthcoming). I have refrained from discussing this strong version of the proposal until now because it seems to me that one could find PA plausible if one does not endorse this sort of mad-dog representationalism.

entirely) because of the difference in their phenomenal character. Perhaps this opens the door once again, for arguing on representationalist grounds that the uber-blindsighter is impossible. That is, one might attempt to argue that since phenomenal characters simply *are* complex representational contents, it is simply inconsistent to maintain that A* and N* can have the same contents and thus are capable of warranting the same beliefs.

I do not find such reductionist representationalist views to be plausible on their own merits. But the debate over representationalism is orthogonal to the problem before us. Representationalism, even of the strengthened variety, does not rule out the possibility of uber-blindsight. To make this point clear, it is useful to consider how representationalists of this strong sort tend to respond to a certain kind of objection to their view. The problem is to explain how experiences in different sensory modalities share contents despite their radical phenomenal differences. Some contents are not shared across modalities: for instance we do not typically hear colors or see the sweetness of sugar. But some contents *do* seem to be cross-modal: for instance properties such as size, roundness, and motion can be represented both visually and tactilely. How is the representationalist to accommodate the view that a subject can be warranted in believing that the object they are feeling is round and that the object that they are seeing is round while the two experiences are radically phenomenally different, and hence must be representationally different?

A standard sort of response is given by Dretske (2003). He points out that in order for such a case to be a counterexample to representationalism, one cannot simply point to two phenomenally different experiences that share *some* content. Instead, one must find a pair of experiences that share *all* contents, but are phenomenally distinct.

This will likely never be true in the cross-modal cases because it will always be plausible that one can get some information from one sense that they do not receive from the other. The phenomenal differences are accounted for holistically, some content is shared, but it is the *total* content of an experience that is identical to the phenomenal character. As a defense of representationalism against the problem of cross-modal contents, this response has some force. However, it is no help whatsoever to PA* in answering the problem we have been discussing. Even if one is a rabid, full-blown, reductionist representationalist, in order to explain the fact that A* and N* warrant some of the same beliefs, one cannot appeal to the intrinsic phenomenal properties of the experiences, one must, again, appeal to the representational contents that the experiences share *despite* their phenomenal differences.

The problem is for the phenomenal approach to specify how A* in the uber-blindsighter's affected field warrants the belief that there is a cube there while N* in her normal field warrants a belief with the same content despite the radical difference in phenomenology. Grant that A* and N* must differ representationally, at least in some ways, in order to accommodate their difference in phenomenology: it still must be the case that there are representational similarities between the two experiences because, despite their differences, they succeed in warranting beliefs with the same contents. This phenomena still calls for explanation, and it is still the case that the explanation will be carried out in terms of features of the experiences rather than their phenomenology.

The lesson of this discussion should be clear: the phenomenological features of one's perceptual experiences simply do not single out all of the epistemically relevant features of one's perceptual experiences. This is so no matter how closely one associates

particular phenomenal properties of experiences with their representational contents. As such, we should reject the phenomenal approach as a plausible answer to the meta-epistemological question.

4. Phenomenal Conservatism

i. What are Seemings?

We now turn our attention to one of the most widely discussed contemporary dogmatist theories: phenomenal conservatism (PC). Originally proposed by Michael Huemer (2001), PC claims that if it seems to one as if p, then one has *prima facie* warrant to believe p. To use Huemer's preferred formulation:

> **Phenomenal Conservatism (PC):** If it seems to S that p, then, in the absence of defeaters, S has warrant to believe p (Huemer 2007).[80]

Huemer means his thesis to be a general epistemic claim, and posits a host of different kinds of seemings: perceptual seemings, introspective seemings, proprioceptive seemings, intellectual seemings, memorial seemings and so on (Ibid.). Only the first of these concerns us here. I will argue that PC is false in the case of perception, and if this argument is successful, then it will follow that the general claim above is false. However, a more targeted version of PC could be the correct epistemological story for another kind of belief or belief-forming process.

[80] In Huemer's formulation he qualifies the view such that seeming that p provides at least *some* warrant for believing p. Huemer claims that this qualification is designed to rule out "weak and wavering appearances." (Huemer 2007) I will not consider such cases here, thus the simplified statement presented above. As always I translate Huemer's talk of justification to my more ecumenical notion of warrant.

Before evaluating PC, let's clarify what the proposal is meant to be. The elements of above definition should be clear given my discussion of these issues so far with one exception: the nature of the *seeming* relation between S and p. Seemings have been much discussed and there is wide agreement on several points in the PC literature. First, seemings are not beliefs or dispositions to believe p. They are not beliefs because it can seem to one as if p, when one does not believe that p. Look at the image below[81]:

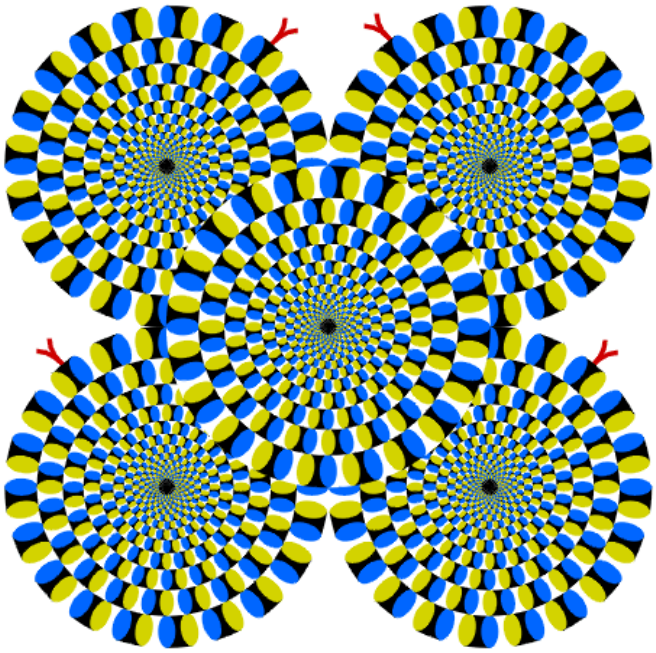

You know that nothing in this image can actually be moving (you are reading this on a piece of paper after all), but it still visually seems to you as if the snakes are moving. Moreover, if you become well acquainted with this illusion it seems plausible that you can lack even any *disposition* to believe that the snakes are moving. You have seen many such illusions before, and you do simply do not have any "tug" or "pull" towards belief at

[81] Image from Kitaoka (2003).

all.[82] Despite this, it will still be the case that it visually seems to you as if p. Therefore, perceptual seemings are identical neither to perceptual beliefs, nor dispositions to believe.

In the case of perceptual seemings this point is particularly plausible. Intuitively, whatever these seemings are, they are likely to be some element of, or identical to, a perceptual experience rather than the resulting belief. Of course, perceptual experiences typically dispose us to believe things, but there can be cases, such as in cases of well-known illusions, where plausibly even the disposition is absent. Citing the considerations just discussed, phenomenal conservatives will often speak as if a seeming is a particular kind of *sui generis* propositional attitude: Cullison (2010, 2013) maintains that seemings are propositional attitudes that are not analyzable in terms of other mental states. Similarly, Huemer (2013a) contends that the concept of "seeming" is unanalyzable and primitive. However, some kind of explanation is in order if we are to come to have any grip on the proposal under discussion.

Seemings, as proponents of PC conceive of them, are a kind of mental state. Perceptual experiences, apparent memories, intuitions and so on are instances of this more general kind. PC claims that mental states of this kind, when undefeated, provide warrant to believe their contents. One might attempt to characterize the view of perceptual warrant that PC recommends as follows: given that perpetual experiences are

[82] It may be that, in some sense I have a disposition to form the belief that the image is moving if I abstract away from other beliefs that I possess. For instance, I may be disposed to form the belief that the snakes are moving were I to believe that the image was an inset GIF or video file on a tablet screen rather than an image on a piece of paper. However, such dispositions, that abstract away from the current beliefs and available evidence of the subject in question, are not relevant to our epistemic evaluations of the subject. Such evaluations are necessarily tied to the evidence, beliefs, and epistemic status of the beliefs that the subject currently possesses. Thanks to Kevin Falvey for calling my attention to these second-order dispositions.

of the general kind "seemings" and seemings warrant beliefs, perceptual experiences warrant beliefs. This characterization, however, would be deceptive. This is because, PC is meant to be compatible with epistemic internalism.[83] In other words, PC must claim that whatever it is that is warranting one's beliefs must be introspectively accessible to the subject. But whether or not a mental state is a *perception* is not introspectively accessible; certain non-perceptual states such as hallucinations and dreams can be introspectively indistinguishable from perceptions. In other words, there are hallucinatory seemings and dream seemings that are introspectively indistinguishable from perceptual seemings, but still (according to PC) warrant belief in their contents.

What follows from these reflections is that if PC is to be an internalist theory, it must maintain that whatever it is about perceptions that typically warrant one's perceptual beliefs is also present in the case in which one forms a belief on the basis of an introspectively identical hallucination or dream. Therefore, defenders of PC typically identify a seeming with some feature of the perceptual experience: its *forcefulness, felt veridicality* or *assertiveness*. The experience is forceful in the sense that it presents its content as reflecting the way the world is. To illustrate the distinction between forceful and non-forceful mental state, consider the difference between a maximally vivid imagining of a visual scene, (so vivid, in fact, that it is phenomenally and representationally identical to a perceptual experience) and a perceptual experience of that same scene.[84] The imagining does not purport to represent the world in the same

[83] Huemer (2001, 2006, 2007), Pyror (2012), Steup (2013) and many others endorse this consequence. Some like Huemer (2006, 2013b) and Steup (2013) even go so far as to claim that PC is necessary to defend internalism from various objections.

[84] This example is due to Huemer (2001).

way as the perceptual experience. Notably, however, deceptive dreams and hallucinations *can* share this feature of perception (at least when they deceive us into believing things).[85] These states, insofar as they are indistinguishable from genuine perception share not just the phenomenology and representational content of perception, but its forcefulness as well.

It is a question of some independent interest how the forcefulness of perceptual experiences is meant to relate to the other two accessible features of experience cited by Huemer: their phenomenology and representational content. Some, like Cullison (2010) and Tollhurst (1998) maintain that the forcefulness is largely phenomenological. They contend that there is something that it is like to have a mental state with this property. Notably, Huemer dissents from this opinion. He presents forcefulness as a distinct feature of experience from representational content and phenomenology (Huemer 2001). Huemer thinks that there are clear cases of seemings without phenomenology. For instance, using proprioception it can seem to one as if one's arm is in a certain location relative to one's torso without any phenomenology associated with it (Huemer 2001, Huemer 2013b).

We need not concern ourselves with these internal debates amongst phenomenal conservatives. All that need concern us here is the characterization that will produce the most defensible version of PC, and we already have strong reason to suppose that defenders of PC shouldn't put much emphasis on any phenomenal character that forcefulness may have. Suppose that there is something that it is like for S to be in,

[85] Strictly speaking, then, PC's claims apply not just to perceptions, but to any *apparent perceptual state* where an apparent perceptual state is one that is internally indistinguishable from a perception as of p. In much of what follows I will ignore this complication and focus on perceptions rather than this broader category.

consider, or be influenced by a forceful mental state. Call the phenomenology associated with forcefulness for S of mental state M, F*. F* will be an instance of what we may call *cognitive phenomenology*. Other examples might be what its like to struggle over a puzzle, the *Eurkea!* feeling one gets when one solves a seemingly intractable problem, what its like to change a deeply held belief, and so on.

The problem for the phenomenally-based version of PC, is that there is no principled reason to suppose that cases of aberrant or altered cognitive phenomenology are impossible. As such we can mount analogous objections to this version of the view as were mounted against the phenomenal approach in the last section. Suppose that what a forceful mental state is like for Will is what it is like for Twill to have a *Eureka!* moment, and vice versa. So when Twill has a state with F* it is not a seeming state, but rather a *Eureka!* moment. Or, if this is difficult to imagine, suppose Will and Twill share no cognitive phenomenology at all; that is there is no phenomenal overlap whatsoever between them. If any such case is possible, then there will be cases in which the a subject is not in F*, but is nonetheless warranted in believing p on the basis of a state with an entirely different cognitive phenomenology. As in the case of the phenomenology associated with particular perceptual experiences discussed in the last section, it is unclear how an appeal to intrinsically different phenomenological experiences can explain the sameness of warrant between the two subjects.

As before, there is a fallback position available to a proponent of a phenomenally-based version of PC. PC, as I have defined it, is a conditional thesis that entails that *if* a subject is in a mental state M with the phenomenal character F*, then the subject is *prima facie* warranted in believing the content of M. Thus, cases of aberrant or missing

phenomenology are not direct counterexamples to such a proposal. The thesis simply does not make claims about the epistemic status of agents who have differing cognitive phenomenology or who lack it entirely. The defender of PC can consistently maintain that PC gives the correct account for agents with F*, but is simply silent about agents who lack this particular phenomenology.

This could be acceptable (though perhaps *ad hoc*) if what we sought was a merely extensionally correct account rather than an explanatory one, but in the present context such a response is unsatisfying. What we want, and what PC was supposed to provide, was an *explanation* of the warranting force of perceptual experiences.[86] But as an explanation of perceptual warrant the current version of PC falls noticeably flat. While it is certainly consistent for a proponent of the view to maintain that both Will and Twill are warranted even though they have distinct phenomenal states, PC, so understood, does not seem to be able to explain *why* both are warranted. Forcefulness, after all, was to be identified with a *particular* phenomenal state F* and Twill is in no such state. Even if a proponent of PC grants that Twill is warranted in his perceptual beliefs, PC itself is silent as to why or how this warrant is generated.[87] Suppose that Twill is warranted in virtue

[86] Perhaps PC was never meant as an explanatory theory at all, but merely attempted to give an extensionally correct account of when we are warranted in believing the contents of our seemings. If so, then the present objection is unfair. Regardless, it is of independent interest whether internally individuated seemings *can* ground an answer to the meta-epistemological question. Furthermore, the Bad Basis Counterexamples of the next sub-section serve to show that PC is not even extensionally correct.

[87] There is an additional problem orthogonal to our current topic, namely that PC, as originally conceived, was meant to be a fully general theory of epistemic warrant, but on the phenomenally-oriented version it only applies to subjects who have the same cognitive phenomenology as us we do. Anyone who endorses PC as a general epistemic thesis, or even thinks it captures the nature of perceptual warrant ought to balk at such a consequence.

of the phenomenal character of his forcefulness-like state F**. If this is the case, PC has no explanation for this warrant, nor for why Will and Twill are epistemically alike in this fashion. That is, they have no explanation for why Twill's seemings warrant his beliefs. I conclude that if we are to appeal to PC to develop an answer to the meta-epistemological question we would be best served by focusing on versions of the view that do not assert that the warranting force of seemings arises from their phenomenology.

ii. Bad-Basis Cases

Michael Huemer does not endorse the version of PC I have just criticized. However, Huemer can be a bit inconsistent in the manner in which he characterizes the property of forcefulness. In some places he quite clearly maintains that the forcefulness is not entailed by the representational content of a mental state, but on the other hand he also characterizes this property of the mental as "representing the content as actualized" (Huemer 2001). We have two options, then: either the forcefulness of a perceptual experience is some additional representational component (i.e. "this very representation is representative of the way the world is"), or it is some distinct unanalyzable introspective feature of the experience. Were I to advocate PC I would prefer the former, as I have a hard time coming to grips with the nature of seemings understood in the *sui generis* fashion Huemer seems to prefer.[88] As it turns out, however, it doesn't matter which of these versions of PC we adopt, because both are subject to the same objections.

[88] Michael Tooley (2013) expresses similar discontent, though I do not accept all of his concerns as serious. For instance, he maintains that the word "assertive" can only apply to people because it is defined in terms of the word "assert." Thus, saying that seemings are "assertive" mental states is, at best, an analogy that requires cashing out in

There are a series of cases that suggest that PC is *overly permissive* in that it ascribes warrant to intuitively unwarranted beliefs. The dialectical ground around the examples I have in mind is well trodden. This being so, I limit myself to a brief description of the relevant cases, the ironing out of some initial problems, and a discussion of why I take a response put forward separately by Huemer and Tucker to be unsatisfying. Markie (2005) presents a series of counterexamples to Huemer's general epistemic thesis of PC. One, however, seems well suited to our current focus on the implications of Huemer's version of PC for perceptual warrant.

> Mental processes that are incapable of producing *prima facie* justified beliefs can nonetheless determine how things seem to us. Suppose that we are prospecting for gold. You have learned to identify a gold nugget on sight, but I have no such knowledge. As water washes out of my pan, we both look at a pebble, which is in fact a gold nugget. My desire to discover gold makes it seem to me as if the pebble is gold; your learned identification skills make it seem that way to you. According to (PC), the belief that it is gold has *prima facie* justification for both of us. Yet certainly, my wishful thinking should not gain my perceptual belief the same positive epistemic status of defeasible justification as your learned identification skills (Markie 2005, 356-357).

In Markie's case, his desire to find gold causes it to visually seem to him as though there were gold in the pan. But the visual seeming is caused in an epistemically improper way. Markie's overwhelming desire to find gold has played an epistemically improper causal role in it seeming to him that there is gold in the pan, and this undermines the warrant

---

more rigorous detail, which is never attempted by proponents of PC. For my part, I think I have an adequate grasp on what an assertive mental state is in the relevant sense, so this problem is of little concern to me. One proposal worth mentioning is that this quality of forcefulness could be meant as a sort of functional notion, that is a representational content has this property insofar as it is, in some sense "treated as true" by the cognitive system of the perceiver. Unfortunately, such functionally-defined characteristics of mental states are not introspectively accessible, and so cannot be made use of in an internalist framework. Furthermore, even if one were to adopt this functionally defined notion of "seeming," thereby ceding the internalist support for the theory, one would still be subject to the counterexamples I survey in what follows.

that his resulting belief that there is gold in the pan enjoys. Furthermore, the case can be constructed so that Markie is not aware of how powerful his own desire for gold is, and thus has no reason to suspect that his desire would influence his perceptual seemings in this manner. It seems perceptually to Markie that there is gold in the pan, and he has no defeaters for this belief. Therefore, according to PC his belief is warranted. Since Markie's belief intuitively lacks warrant, the case presents a counterexample to PC.

One might be inclined to point out that the belief that there is gold in the pan is not perceptually basic. The perceptual system likely cannot represent chemical kinds such as "gold," and so one might contend that the PC proponent is not committed to saying anything about such a case. First, it should be noted that Huemer would accept no such restriction thinking, as he does, that *any* seeming, perceptual or not, provides *prima facie* warrant to believe its content. However, one may attempt to salvage some version of PC by limiting its scope, as other dogmatists are inclined to do, to basic perceptual beliefs. However, nothing about the case seems to depend on the relevant belief being about a chemical kind like *gold*. McGrath (2013) presents a revised version of Markie's case in which the pebble in the pan is really only a dullish yellow. To the expertly trained observer, or indeed anyone without an overwhelming desire for gold it does not appear to be the color of gold, or glitter. However, the greedy subject's desire for gold is so great that to him the pebble seems to be the paradigmatic color of gold and to glitter in the sun. If such a case is possible, then even a restriction to basic perceptual contents, of which color and glittering are plausible candidate, will not remove the problem PC faces.

Furthermore, our background beliefs and/or past experiences can alter, in surprisingly fundamental ways, the contents of our perceptual experiences. Hansen et. al.

(2006) conducted an experiment in which the subjects were asked to adjust the color of an image of an object with a paradigmatic color, for instance a banana, until the image appeared to them to be achromatic. Surprisingly, the subjects would change the color of the image of the banana until it was a slightly bluish tint indicating that when the banana was *genuinely* achromatic, it still appeared to them to be yellowish. Control images of random dots, or shapes without paradigmatic colors produced no such effect. This is held to be evidence of *cognitive penetration* of background beliefs into the content of the subject's perceptual experiences.

Presumably the subjects of Hansen's trial were warranted in believing that bananas are typically yellow, and so it may be thought that such a case does not impugn their warrant for believing (falsely) on the basis of their perceptual experience that the achromatic image of the banana was slightly yellowish. However, we need only change the case in order to produce a problematic case for PC. Suppose that one of the subjects of Hansen's trial, Sandra, was from a remote island, had never seen a banana or a color-picture of the fruit but has seen many black and white images of bananas. Furthermore, she has not been informed by any reliable source about their ordinary color. Sandra is, however, a devout believer in the practice of divination, though she has no reason to believe in the predictive reliability of the practice, and indeed knows that the practice has led her to many false beliefs in the past. Earlier that day, elbow deep in squirrel entrails, she came to believe that bananas are, in fact yellow. So deep is her commitment to the veracity of divination, that she comes to believe this fact as strongly and completely as

anyone who had been around bananas their entire life. When she takes Hansen's test later that day the achromatic image of the banana appears to her to yellowish.[89]

Is Sandra warranted in believing that the achromatic image is yellowish? It seems not. She has the experience that she has partly because of her thoroughly unwarranted belief that bananas are typically yellow. There seems to be a significant epistemic difference between her belief and that of the other subjects in the trial. *Their* experience of a yellowish image occurs partly in virtue of their warranted background beliefs about the typical color of bananas, whereas Sandra's results from an entirely unwarranted belief with the same content. However, once again, PC predicts that there is no epistemic difference between Sandra and the other subjects. Hence, PC is false. Call these **Bad Basis Cases.**[90] Such cases seem to indicate that if a perceptual experience has the content that it does because it is formed on the basis unreliable or epistemically blameworthy processes, then the resulting perceptual belief will not be warranted. PC, however, claims that *all* perceptual seemings provide *prima facie* warrant for their contents, and so cannot distinguish genuinely warranting instances from their Bad Basis counterparts.

[89] One may be concerned that simply having a strong belief would not be sufficient to produce the effect noted in Hansen et. al.'s experiment. Rather, one might think that it is something about having many perceptual experiences as of yellow bananas that produces the effect. To accommodate this worry, we could stipulate that Sandra carries out her divination ritual under the influence of a powerful hallucinogen causing a hallucinatory state where she experiences a rapid succession of yellow-banana images.

[90] DePoe (2011) repurposes BonJour's (1985) clairvoyance case as a similar sort of counterexample to PC. Suppose that Norman's unreflective clairvoyant ability causes it to seem to him that Barack Obama is in the Oval Office, and thus Norman believes that Obama is in the Oval Office. DePoe contends (and I tend to agree) that it plausible that in virtue of this seeming Norman is warranted in his belief in Obama's location. Huemer (2011) gives essentially the same response to this case, as his response to Case 3 that I discuss below.

iii. A Response on Behalf of Phenomenal Conservatism

In response to such cases, Huemer (2013b) begins by distinguishing between three different sorts of Bad Basis Cases. Let M be the mental state that improperly causes the perceptual experience that p:

> Case 1: The subject is aware that M caused her to have the perceptual experience, and the subject has evidence that an experience caused by M is unreliable. (Huemer 2013b, 343)

As Huemer correctly points out, such a subject has a defeater for her belief that p that undermines the *prima facie* warrant the experience provides. Therefore PC is not committed to claiming that she is justified in believing p. However, none of the cases described above were of this form. Rather, in each it was explicitly stated that the subject was unaware of the epistemically problematic nature of the background mental states that improperly caused her experience. This case, therefore, is irrelevant to the current dispute. Let us move on:

> Case 2: The subject is aware that her experience was in some way caused by M, but neither doubts nor has reason to doubt that M is a reliable source. (Ibid. 344)

Huemer acknowledges the difficulty of imagining such a case, but contends that insofar as it is possible, he sees no problem with attributing warrant to the subject. He maintains (Huemer 2013a, 344) that if one were aware that one's belief were aberrantly caused, this would not undermine the warrant unless one had some positive reason to think the aberrant causation introduced an element of unreliability. He contends that there is no counterexample to PC in such a case. Once again, the cases above do not have this feature. None of the subjects in the trial were aware of the effects of cognitive

penetration that Hansen et. al. (2006) were testing for, and in the gold-digger test it was simply stipulated that greedy prospector was not aware of his own intense desire, and thus had no reason to suspect it would be influencing his experiences.

Still, Huemer's claims regarding Case 2 are highly suspect. In particular, his understanding of what it takes in order for a belief to be defeated seems implausibly strict. Huemer is guilty of the error that Wright identified and criticized in the argument I discussed in Chapter II Section 3 of this dissertation. Recall that Wright (2007) claimed that the dogmatist must understand "reason to doubt" as having positive evidence against the veridicality or reliability of one's perceptions. Understood this way, Huemer can respond to Case 2 as he does above. However, the view is then subject to counterexamples of the kind Wright constructed. Suppose that I find out that my present perceptual experience was caused by Martians shooting a beam into my brain. Suppose, also, that I have no positive reason to think that they would give me a non-veridical experience because I know nothing about their motivations. For all I know, they are running a test, and the experience that they implanted is precisely the one I would have without their interference. Indeed, for all I know, they are *correcting* non-veridical elements of my perception. Still, based solely on my evidence that my perceptual experience was aberrantly caused, my perceptual warrant is defeated. It seems that having evidence of aberrant or unusual causation is enough (absent additional evidence that the aberrant cause is not equally or more reliable than the ordinary one) to undermine my perceptual warrant. Thus, one should say, of the subject in Case 2, that she is unwarranted in her perceptual belief. In response to such examples I argued that the dogmatist need not, and ought not, endorse Wright's claims about the nature of defeat. In

order to avoid objections of the kind posed by Wright, therefore, a proponent of PC ought not endorse Huemer's claims concerning Case 2.

However, this point, while problematic for Huemer in particular, is not a serious problem for PC itself. If one accepts the amended understanding of defeat advocated here and in the previous chapter, then one could simply state PC making use of this notion. This supplemented version of PC would then give what I take to be the right result with regards to Case 2, namely that the subject is unwarranted. Moreover, this is entirely consistent with internalism because it is based on evidence of aberrant causation that the subject is aware of. So revised, Case 2 would present no counterexample to PC. Thus, I agree with Huemer that cases of this form do not constitute counterexamples to his view, though he is incorrect about why this is so.

The last variety of Bad Basis Case Huemer discusses have the form of those I presented in the previous sub-section:

> Case 3: They subject is unaware and has no evidence suggesting that the appearance is caused by M.

Huemer maintains in this case that there is no clear counterexample to PC "because the subject actually *would* have some justification for believing p. When the subject is unaware of an appearance's etiology, that etiology is irrelevant to what it is rational for the subject to believe" (Ibid. 344). Given that all of the putative counterexamples were of just this form, this response amounts to simply biting the bullet. If one is already convinced of the truth of PC, this may be a satisfying response. However, Huemer's response will be unpersuasive to any who, like me, found it problematic to attribute warrant to the subjects in the Bad Basis Cases described above.

Huemer (2013b) and Tucker (2010) attempt alleviate residual contrary intuitions by distinguishing between the conditions necessary for perceptual *warrant* and those necessary for perceptual *knowledge*.[91] The greedy gold-digger, they claim, is *warranted* in believing that the pebble is golden because to believe this is the correct response to his improperly caused perceptual experience. However, his belief is not a candidate for knowledge in virtue of the fact that it is improperly caused. Similar remarks would apply to Sandra's case of cognitive penetration of her perceptual experience by unwarranted beliefs. PC is held up as an adequate account of warrant while it is granted that in order for one to have knowledge additional, non-introspectively accessible conditions must be met.

According to this suggestion, features of one's belief forming methods that externalists focus on: e.g. a proper etiology, the proper functioning of one's perceptual system, and reliable veridicality, are conceded to be conditions for knowing some proposition, but are rejected as conditions for warrant. Nearly everyone, post-Gettier (1963), will allow that warranted true belief is not sufficient for knowledge. Huemer and Tucker seem to be appealing to a distinction that nearly everyone would grant in order to explain why their judgments about the Bad Basis Cases are not as implausible as they initially seem. They are not committed to claiming that one can come to *know* something on the basis of improperly caused or epistemically suspect experiences, only that one can be warranted in one's beliefs based on them. Bad Basis Cases, then, simply constitute

[91] Huemer makes the point by utilizing Plantinga's notion of warrant according to which "warrant" is whatever must be added to true belief in order to secure knowledge. In order to avoid terminological confusion, I use Tucker's formulation of the response.

another variety of Gettier-like cases in which one's belief can be warranted and true but not amount to knowledge.

iv. Norms and Warrant

I do not find the preceding response on behalf of PC satisfying. To see why, let's get back to epistemological basics. There are several ways in which a belief can fail with regards to the goal of securing knowledge. The clearest and simplest way for it to do so is to be false. However, simply failing to be true does not entail that a belief was formed improperly, in light of bad evidence, or is in any other way deficient. Sometimes the evidence that one has is deceiving, or the environment that one finds oneself in is uncooperative such that using otherwise excellent belief-forming methods, one may end up with a false belief. In such a case, the belief is normatively above reproach, even though false. Moreover, Gettier (1963) and the legions of commenters in his wake have pointed out that there are many cases in which a particular sort of mismatch between one's warrant, and the nature of one's environment can result in failures of knowledge even if the belief is epistemically above reproach, and true. The world can fail to cooperate in more complicated ways than simply rendering one's belief false, without impacting the warrant the belief has. It is a matter of widely accepted philosophical dogma that one's belief can be warranted, but fail to amount to knowledge even though it is normatively above reproach.

However, in the Bad Basis Cases above the reason the subjects' beliefs are not suitable for knowledge is precisely because they have been formed in a normatively deficient fashion. The greedy gold-digger forms his belief on the basis of a kind of

wishful thinking in which his desires have exerted an impermissible effect on his experiences and, via his experiences, his beliefs. Similarly, our divination fan's belief that the grey banana image was yellowish was formed in part because of her deeply irrational beliefs in the power of divination generally, and the testimony of squirrel intestines in the particular case of banana colors. The failures of these agents are not of the same kind as one typically encounters with Gettier cases, cases of warranted false beliefs, or other cases in which the world simply fails to cooperate with a subject's perfectly acceptable epistemic practices. Rather, the subject's failure to have knowledge in the Bad Basis Cases depend crucially on the epistemic failings of the subjects themselves. It won't do, then, to maintain that the subjects in these cases fulfill the normative conditions on knowledge but fail in some further way that does not implicate the epistemic grounding of their beliefs. Rather than some non-normative mismatch between the world and their beliefs, the subjects are epistemically deficient and fail to have knowledge precisely because they lack warrant for their beliefs.

This is not to deny that there are ways of evaluating the subjects in the Bad Basis Cases according to which they believe what they ought to believe. For instance one can wonder, abstracting away from particular features of the cases, what one ought to believe when one is confronted with a particular perceptual experience. One can reasonably ask what, given a subject's evidence, experiences and background beliefs at a particular time, they ought to believe abstracting away from the epistemic bona fides of the background beliefs and the experience itself. To illustrate this sort of evaluation, consider a case of deductive reasoning from unwarranted premises. Suppose that a subject has a set of entirely unwarranted beliefs, p, q, and r. Suppose that these beliefs collectively entail a

further proposition s, and the subject knows this. There is a clear sense in which the subject ought to believe s. In holding unwarranted beliefs in the first place, the subject fails epistemically, but given that the subject holds the beliefs that she does, and knows about the entailment, it would amount to an *additional* epistemic failure if she were to fail to believe s. None of this should suggest, however, that the subject's belief in s would be warranted on the basis of the inference, or that she could come to know s in this manner if s is true.

A similar diagnosis is appropriate in the Bad Basis Cases. It may very well be true that, given that Sandra has an experience as of a yellowish banana image, she ought to believe that it is yellowish in the sense just articulated. Whatever her other epistemic failings, to refuse to form a belief that an object is the way that it visually appears to her to be would amount to an additional epistemic failure to appreciate the (usual) warranting power of perceptual experiences. However, this should not incline us towards the conclusion that Huemer and Tucker advocate that Sandra is, thereby, warranted in her perceptual beliefs. PC claims that having an experience with the particular feature *forcefulness* is sufficient to secure warrant. However, background cognitive facts of which a subject is herself is unaware can play a significant role in determining whether or not a particular experience is epistemically respectable. If perceptual experiences can be cognitively penetrated by one's beliefs and past experiences or by one's desires in the manner described in the Bad Basis Cases, this interference may result in the experience failing to confer warrant.

5. Reconsidering the Internalist Intuition

In this chapter I have argued that none of the introspectively accessible elements of perceptual experiences can ground a satisfactory answer to the meta-epistemological question. I will close this chapter by considering once again the motivations for an internalist approach to the meta-epistemological problem. Recall Huemer's schema for generating problematic cases for externalism:

1. S satisfies E with respect to p, but not to q.
2. S satisfies all other parts of C with respect to both p and q.
3. S satisfies no other sufficient conditions for warranted belief in p or q.
4. It seems to S as if he is in the same epistemic position with respect to p as he is to q, and he has no reason to suspect that one is more warranted or likely to be true than the other.
5. S ought (epistemically speaking) to believe p and not believe q.

E is meant to be some externalist condition or other for the possession of warrant and C is the set of all necessary and sufficient conditions for a warrant of the belief in question. This posed a dilemma to the externalist: continue to maintain that E is a genuine condition for warranted belief, in which case what one ought to believe can come apart from whether or not one is warranted, or to conclude that since p and q are epistemically on par for S, he ought to believe and is warranted in believing both propositions, hence E is not a condition required for warrant after all.

This reasoning relies on a strong connection between what one ought, epistemically speaking, to believe, and what one is warranted in believing. In Section 2, I granted that it is difficult to see how it could be that one ought, epistemically, to form a belief that is unwarranted. If the reflections of the last section are correct, however, we must be wary of an equivocation between different dimensions of epistemic evaluation. It is quite possible, as in the case of considering the deductive entailments of one's

unwarranted beliefs, for it to be the case that one *ought* to believe some proposition, while it is quite clear that even if one does, one will lack warrant for that belief. If one takes a narrow view on the case and is interested only in what the subject ought to do given her current doxastic state, then it is quite clear that failing to believe known entailments of her beliefs constitutes an epistemic failure that she ought to avoid. This is the sense in which she ought to believe the logical consequences of her beliefs. Thus, it is possible for one to lack warrant for believing some proposition while it still being the case that (in some sense) one ought to believe it.

This suggests that Huemer's schema will be of little use in drawing conclusions about the nature of perceptual warrant unless one can be assured that the "ought" in the final clause singles out the evaluative norms specifically associated with *warrant* rather than any of a host of other dimensions of epistemic evaluation, for instance: rationality, epistemic blamelessness, or coherence. The nature of the norms associated with each of these is of independent interest and each deserves investigation. Such questions are of enormous importance generally and of specific importance for certain sorts of belief formation, (for instance: rationality and coherence in particular would seem to be very significant for understanding inferential reasoning as well as understanding the epistemology of theory building as in the sciences, and in folk theories like folk psychology, and folk physics, and so on). However, recognizing the potential equivocation over epistemic "oughts" renders the general schema Huemer presents largely useless. If a putative counterexample to some externalist conditions on warrant that satisfies the schema is offered, we must examine the details of the particular case in order to ascertain whether or not the manner in the sense in which the subject "ought to

believe" is best cashed out in terms of the norms associated with *warrant* rather than some other dimension of epistemic evaluation. It is open, then, to the externalist to seize the first horn of the internalist's dilemma and maintain that it is not nearly so problematic as one might initially think to maintain that a subject ought to form an unwarranted belief, so long as one is clear about the different dimensions of epistemic evaluation in play.

In response to Huemer's schematic objection, I have given a schematic response. Of course, it may be possible for one to construct a case satisfying Huemer's schema in which no such move is plausible. One may think that I have already done so in the case of my BIV duplicate, BIV-Tim. I will, however, put off discussion of this particular problem until the final chapter of this dissertation. The reason for my reticence is that I take its force to depend crucially on what sort of externalist view one advocates. For instance, depending on the form of externalist dogmatism that one advocates, it will be true, at least in some BIV cases, that BIV-Tim holds warranted (though false) beliefs. In particular, I think that one substantial benefit of the externalist version of dogmatism that I favor is its ability to give what I take to be a satisfactory response to this very problem. I will return to this issue in Chapter IV. Given the failures of internalist responses to the meta-epistemological question, I will proceed to develop an externalist version of dogmatism according to which one's immediate perceptual warrant is to be explained by one's *perceptual competence*.

Chapter IV. Competent dogmatism

1. Introduction

In this chapter I will develop a positive theory of perceptual warrant. According to the theory that I defend, one's basic perceptual beliefs are immediately and defeasibly warranted if they are formed on the basis of experiences produced by a *competent perceptual system*. Perceptual systems produce representations in the form of perceptual experiences. Representations entail norms associated with veridical representation. That is, a representation of the world can fail or succeed with regards to whether it represents its subject matter veridically. A perceptual system (for instance a visual system) is, therefore, subject to norms of veridicality simply in virtue of the fact that its output experiences represent features and properties of the local environment of the organism. This does not entail that veridical representation must be an etiological function or aim of the visual system. It doesn't have to be the case that reliably veridical representation was evolutionarily selected for, or that the primary biological/evolutionary function of the visual system is to produce veridical representations. It is possible that a subject could survive and thrive without a cognitive system that reliably represents its environment. It may be that representing the environment unreliably is conducive to the survival of an organism. While such perceptual systems succeed relative to the goal of promoting the survival of an organism, they fail with regards to the goal of veridical representation. An ability to succeed relative to the goal of survival does not entail that the subject succeeds qua veridicality.

I will present a view according to which it is the veridicality-aimed norms that are relevant to the epistemic evaluation of a perceptual system and a perceiving subject. The

degree to which a perceptual experience warrants a subject in believing its contents will be tied to the degree to which the perceptual system that generated the experience succeeds in representing the world. I will argue that features of one's perceptual systems (that are not introspectively accessible) can render them competent to represent a particular environment. When these conditions are met, a subject is warranted in forming beliefs on the basis of her perceptual experiences. In a slogan: if one has a perceptual experience as of p, and that perceptual experience was produced by a competent perceptual system, one has immediate defeasible warrant to believe that p. I call this claim *competent dogmatism*.

2. Competence and Reliability

I will start by discussing the reasons that the reliability of one's perceptual systems does not suffice for perceptual warrant. I will argue that the failure of reliabilism introduces two important constraints on a discussion of perceptual warrant. First, a perceptual system may be reliable with regard to one's environment, and unreliable with regard to others. What is required is a non-arbitrary way in which to fix the epistemically relevant environment of evaluation for perceptual systems. Second, a perceptual system (indeed any belief-forming process) may be reliable for accidental, lucky, or highly contingent reasons. In such cases it seems implausible that the subject's perceptual system is competent in a way that would allow to support warranted perceptual beliefs for that subject. An adequate theory of perceptual warrant should provide principled distinctions between such cases and cases in which subject's are intuitively warranted in their perceptual beliefs.

Goldman (1979) proposed the following general theory of warrant:[92]

> S's belief in *p* at time *t* is warranted just in case it meets one of the following two conditions:
>
> (1) S's belief in *p* at *t* results from a belief independent process that is unconditionally reliable.
> (2) S's belief in *p* at *t* results from a belief-dependent process that is conditionally reliable, and the beliefs on which the process operates are themselves warranted (Goldman 1967, 117).

In Goldman's terminology, perception is a belief-independent process because it does not take beliefs as inputs to its processing.[93] Goldman's theory entails, therefore, that if forming beliefs on the basis of one's perceptual experiences is reliable, that is yields (and in nearby possible worlds, would yield) more true beliefs than false ones,[94] then one can form warranted beliefs on this basis.

An externalist dogmatist could contend that this provides an initially satisfying answer to the meta-epistemological question of Chapter III. Perceptual experiences are capable of warranting beliefs because they are good guides to the truth. If one forms beliefs on the basis of one's perceptual experiences, one is more likely to end up with true beliefs, therefore one's basic perceptual beliefs will be warranted just in case one is

[92] As has been my custom, I reformulate Goldman's terminology to match my own.

[93] If one wishes to maintain that at least some cases of perception (e.g. expert perception, learned associations, etc.) do take beliefs as inputs, this claim could be restricted to the basic perceptual beliefs that the dogmatist is most directly concerned with.

[94] How high the ratio must be could be a matter of some debate, and is probably vague. Goldman freely acknowledges as much (Goldman 1979), and takes it to be an advantage of his theory. Warrant, just like reliability, comes in degrees. A more careful formulation of the view would have it that a belief is warranted to the degree that the process that produced it is reliable. Warrant *tout court*, that is warrant sufficient for knowledge, requires some (probably vague, and perhaps context dependent) degree of reliability of the belief-forming process.

in possession of a sufficiently reliable perceptual system. Thus, if the general reliabilist epistemology presented above is correct, the dogmatist will have a ready explanation for why one's experiences can suffice to warrant one's beliefs. However, for well-known reasons this account will not do. In the rest of this section I will discuss some of the most frequently cited objections to reliabilism. These objections are important because the failures of reliabilism present desiderata that an adequate externalist dogmatist theory must fulfill.

i. The Domain Problem

One problem for reliabilism has to do with specifying the epistemically relevant domain of evaluation for a belief-forming process. A belief-forming process may be reliable in one environment, and unreliable in another. As we will see in subsequent sections, the reliability of human perceptual systems requires that the system encode assumptions about the contingent features of the subject's immediate environment. If placed in an environment without these features, the perceptual system may fail to produce reliably veridical perceptual experiences. However, the fact that a perceptual system fails to reliably deliver true beliefs in some distant possible environment should not entail that the organism, in its ordinary environment, cannot form warranted beliefs on the basis of its perceptual experiences. The organism's normal environment is the one in which it must ordinarily navigate, hunt, interact with, and survive in. Distant environments that bear no relation to the organism's ordinary aims and practices, and no causal relation to the way in which it ordinarily forms beliefs, are inappropriate domains for evaluating its system in environments in which it ordinarily functions. The problem for the reliabilist is

to specify, in a non-*ad-hoc* and explanatorily satisfying manner, what domain or environment is the epistemically relevant one for evaluating the reliability of a particular belief-forming process. Call this **the Domain Problem**.

A simple solution to the Domain Problem is to evaluate the reliability of the process with regards to the environment in which the subject formed the belief. This account has the virtue of dealing with the case considered above. According to this proposal, we need not consider how its perceptual systems would function on some distant planet or in another possible world. All that matters is whether its perceptual systems reliably produces veridical representations in the environment in which it presently finds itself. Unfortunately, this simple solution has counterintuitive consequences. These consequences are made clear by considering what has come to be known as the New Evil Demon Problem.[95] Consider my brain-in-a-vat counterpart, BIV-Tim. BIV-Tim, recall, is a human brain that has been placed in a vat connected to a super-computer which stimulates the sensory parts of his brain to have the same perceptual experiences as I do. BIV-Tim has all of the same experiences that I do, has all the same evidence available to him that I do, and forms the same beliefs that I do on the basis of these experiences and evidence.[96] I am a normal embodied human who typically

[95] I introduce the problem by its more common name. Due to various questions about the metaphysical possibility and/or conceivability of Cartesian Demon cases, I carry out my discussion of the problem in terms of a different radical skeptical scenario: that of a BIV.

[96] One may contend, as does Putnam (1981) that certain versions of such cases are inconsistent with content externalism, a thesis I have been at pains to remain consistent with throughout this dissertation. I will discuss a range of BIV cases and their implications for the view I develop in Section 8. I put aside concerns about BIV cases and their possibility until then, and proceed with the simplified and intuitive description of the case presented above.

interacts with physical objects that are largely as I believe them to be. BIV-Tim, on the other hand, is a brain-in-a-vat who is receiving deceptive experiences as of physical objects, even though he never interacts with such things at all. BIV-Tim is radically (though blamelessly) deceived about his environment. However, despite his being radically deceived, it does not seem correct to claim that BIV-Tim is less warranted in his beliefs than I am. After all, we have formed our beliefs for precisely the same reasons. If BIV-Tim and I were to switch places, we would both continue to hold the beliefs that we have held, and to form and revise our beliefs using the processes that we have used up until the time of the switch. Since we form our beliefs using the same processes, for the same reasons, and on the basis of the same evidence, it is very plausible that if *I* am warranted in my beliefs, it must be the case that BIV-Tim is warranted as well.

Considering this case presents two problems for the answer to the Domain Problem we are considering. First, because my belief-forming processes are reliable and BIV-Tim's are not, reliabilism entails that I am warranted in my beliefs while BIV-Tim is not. If one takes the domain of evaluation for belief-forming processes to be the environment in which a subject actually finds herself, this means that one cannot respect or explain the intuitive epistemic similarity between BIV counterparts. The second problem has to do with what happens when BIV-Tim and I switch places. Suppose that the scientists responsible kidnap me one night and replace my brain with that of BIV-Tim and hook me up to his vat. Up until the moment of the swap, according to reliabilism I am perfectly warranted in the vast majority of my beliefs. However, when I am transported (by completely imperceptible means) to the vat, reliabilism entails that I can no longer form warranted beliefs about my immediate environment. This is because none

of my perceptual systems are reliable in the new environment. Like BIV-Tim before me, I am now radically deceived. But this is counterintuitive. I have not changed my processes for belief formation, nor have I received any defeating evidence that would undermine my trust in those processes. Simply moving a subject to a different environment should not result in such a radical shift in its ability to form warranted beliefs.

The New Evil Demon Problem is sometimes taken to constitute a refutation of the claim that reliability is a necessary condition on warranted belief-formation (Smithies 2014). It should be clear from my presentation that no such argument will be sound. Even granting the intuitive judgments concerning the case I endorse above, at most what it shows is that the warrant-relevant domain of evaluation for belief-forming processes is not the environment in which the subject currently finds herself. Reliability in some epistemically privileged environment may, for all that has been said, still be a necessary condition for the formation of warranted perceptual beliefs. This does, however, present a difficult problem for any epistemic theory that endorses the claim that reliability is a necessary condition on warranted belief-formation; namely, to specify in a principled way which is the relevant environment for evaluating the reliability of a subject's belief-forming processes. A satisfying solution to the Domain Problem must not only select an environment of evaluation that yields intuitively correct verdicts about which of a subject's beliefs are warranted, it must also do so in a principled and explanatory manner.

To see the difficulty of meeting the second of these desiderata consider Goldman's (1986) attempt to provide a solution to the Domain Problem. That Goldman foresaw this class of objections in his initial presentation of reliabilism is clear from his

discussion of counterfactual worlds in which wishful thinking is reliable due to the interventions of a benevolent and omnipotent demon (Goldman 1967, 120-121). Goldman initially thought to deal with this difficulty by stipulating a class of "natural" or "unmanipulated" worlds as the epistemically relevant domain of evaluation (Ibid.). In later work, he spelled out this general idea in more detail: claiming that one should evaluate the reliability belief-forming processes with regards to *normal worlds* (Goldman 1986). A normal world is a world "consistent with our general beliefs about the world" (Goldman 1986, 107).[97] According to this new proposal, it doesn't matter that BIV-Tim's belief-forming processes are radically unreliable in his world, because his world is not *normal*. Since BIV-Tim would, by stipulation, form the same beliefs in our world as I do, and since *my* belief-forming processes are reliable in this world, BIV-Tim's belief-forming processes are reliable in this world as well. Since our world is a normal world according to Goldman, BIV-Tim's belief-forming processes are reliable in normal worlds and therefore BIV-Tim's beliefs are warranted, even if they tend to be false in the world in which he lives.

[97] There are a host of difficulties concealed in this formulation. First, whose beliefs matter? Goldman uses a general pronoun "our" but it is unclear what the domain is meant to be. Is it meant to apply to all people? All sane people? All people that share our general scientific worldview? Second, which of our collective beliefs are the one's that are general enough to pick out a normal world? That the world is physical, and that it is not a simulation are surely candidates, but what else gets into this class of beliefs? Finally, if we select out a class of beliefs, it could be the case (if they are general enough) that many worlds that are very different from one another meet these criteria. It could be the case that our belief-forming processes are reliable in some of these worlds, but not others. It is unclear, on the present proposal, how we are to evaluate the reliability of our belief-forming processes if this is the case. I don't hold out much hope for satisfying answers to these questions. Nor does Goldman (1992, 136). I do not emphasize these issues here because I take it that there is a more fundamental problem that this account faces even if these issues could be adequately addressed.

Goldman later acknowledged this proposal to be a failure (Goldman 1988). Normal worlds, on his account, are picked out relative to *our* (that is normal humans') general beliefs about the world. This introduces an implausible chauvinism into our epistemic judgments. Consider a world sufficiently different from ours so that it is not a normal world. Suppose that in this world there are subjects whose belief-forming processes are adapted to the peculiar local conditions in such a way that they reliably produce true beliefs. Suppose further that their belief-forming processes are unreliable in worlds like ours. According to Goldman's (1988) account, these subjects cannot form warranted beliefs because they would not reliably do so in a normal world because, by stipulation normal worlds are *worlds like we believe ours to be*. This consequence is untenable. Why should it matter if such alien subjects are unreliable in "normal worlds"? Surely they ought not form beliefs as if live in (what we would describe as a) "normal" world, as this would render most of their beliefs false and would frustrate those of their projects that are predicated on their perceptual beliefs. Rather, they ought to use processes that produce true beliefs at a high ratio in *worlds like theirs*. Clearly we need another way to individuate the correct environment of epistemic evaluation.

There is much more that could be said about the Domain Problem and how a reliabilist might seek to address it. I will return to these issues in Sections 5 and 8. At this point, I would only reiterate that the Domain Problem is not only a problem for reliabilism. It is a problem for *any* epistemological theory that takes reliability to be a necessary condition on warrant. In particular, if one takes the reliability of one's perceptual system to play some important role in explaining its ability to produce warranted beliefs, one must have a non-*ad-hoc* and explanatorily satisfying way of

specifying the environment that is relevant for the evaluation of the reliability of the belief-forming method in question.

ii. The Contingency Problem

BonJour (1985) and Lehrer (1992) have presented two cases that show that a reliable belief-forming process can fail to provide even *prima facie* warrant for one's beliefs.

> **Norman the Clairvoyant (BonJour 1985)**
>
> Norman, under certain conditions which usually obtain, is a completely reliable clairvoyant with respect to the location of major American political figures. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power, or for the thesis that he possesses it. One day Norman comes to believe that President Obama is in New York City, although he has no evidence pertaining to this belief. In fact, the belief is true and results from his clairvoyant power under circumstances in which it is completely reliable.

> **Mr. Truetemp (Lehrer 1990)**
>
> Mr. Truetemp has unknowingly had a device surgically implanted in his head. The device is connected to a satellite that feeds it highly accurate data about the weather outside of Truetemp's house. Upon receiving information that the temperature outside is X the device electrically stimulates Truetemp's brain in such a way that he forms the belief that the temperature outside is X. Truetemp has no evidence for the existence of this device, and he has no evidence that his spontaneously formed beliefs about the temperature outside are reliably true. At noon on a Friday, Truetemp is eating lunch in his air-conditioned windowless kitchen. He has no antecedent evidence about the temperature outside. The device in his head receives a signal that the temperature outside is 74 degrees and stimulates Truetemp's brain in such a way that he forms the belief that the temperature is 74 degrees outside.

Norman's belief about the location of the president and Truetemp's belief about the temperature outside of his house are intuitively unwarranted in the cases as presented. BonJour and Lehrer took these cases to show not only that reliability does not suffice for justification, but also that warrant requires meta-beliefs about the reliability of one's

belief-forming processes. Lacking such meta-beliefs, they contended, neither Norman nor Truetemp are warranted in their beliefs.

A major theme of this dissertation has been that demanding meta-beliefs of this kind over-intellectualizes the conditions for perceptual warrant. A subject need not possess meta-beliefs about the reliability of her visual system in order to form warranted beliefs on the basis of her visual experiences. Furthermore, this over-intellectualized constraint is not entailed by our intuitive judgments about the Norman and Truetemp cases. That Norman and Truetemp lack meta-beliefs about whether their beliefs are likely to be true is hardly the only epistemically significant feature of these cases. Neither Norman nor Truetemp have experiences or other antecedent mental state on which their beliefs are grounded. Their abilities are due to artificial or supernatural features beyond the normal functioning of ordinary humans; they did not acquire these abilities either by interacting with their environment, or by inheriting them from their predecessors. In Norman's case, there is no specified physical mechanism that explains how he acquires information about the president's location and we are left to imagine some magical, mysterious connection. Truetemp's abilities have very little to do with the competence or reliability of his own cognitive mechanisms. His abilities are artificial in a way that doesn't seem strongly connected to an evaluation of Truetemp *himself* as an epistemic subject.

All of these features could be appealed to in order to explain our intuitions about Norman and Truetemp. To illustrate this point, consider the following revised cases due to Lyons (2009):

Nyrmoon

Nyrmoon is a member of an alien species for whom clairvoyance is a normal cognitive capacity, which developed in much the same way as vision does for humans. Members of Nyrmoon's species have specialized internal organs that are receptive to the highly attenuated energy signals from distant events; as an infant, all was a "blooming buzzing confusion" for Nyrmoon until, like everyone else, he learned to attend selectively, recognize objects, and filter out coherent distant events. Nyrmoon, however, is so extremely unreflective that he has no beliefs about the reliability of his clairvoyance. One day he forms, as a result of clairvoyance, the belief that his house is on fire (Lyons 2009, 119).

Mr. Vipertemp

Mr. Vipertemp is a member of an alien species that has, like our own pit vipers, evolved a sensitive and highly reliable heat-detection faculty, complete with special thermoception organs. If Vipertemp is sufficiently unreflective, he may lack any specific meta-beliefs about the reliability of his thermoception. Using this thermoception, he forms the belief that it is 104 degrees outside (Ibid).

Nyrmoon and Vipertemp are intuitively warranted in their beliefs. In these cases, their warrant is explained by their possession of perceptual systems that we lack. That a subject has different perceptual capabilities than normal human beings does not suffice to undermine that subject's warrant. These revised cases indicate that we should not take BonJour's and Lehrer's original cases to establish the need for some highly intellectualized constraint on perceptual warrant (Ibid.).

On the other hand, the asymmetry between the original and revised cases do seem to support the contention that reliability, while perhaps necessary, does not suffice for warranted belief-formation. Thus, they do present counterexamples to Goldman's classical version of reliabilism. Goldman himself introduced a case of this kind when considering a benevolent all-powerful demon that uses his omnipotence to make it the case that a certain intuitively non-warranting belief-forming process reliably produces true beliefs. In his initial formulation, the demon makes it the case that all of a subject's

beliefs formed by wishful thinking are true (Goldman 1967, 120). It seems that even though wishful thinking is reliable for this subject in the benevolent demon world that she is not warranted in the beliefs she forms on its basis.[98]

We can construct a similar case that specifically concerns the set of belief-forming processes that are the focus of this dissertation, namely beliefs formed via perception. Consider a subject with a perceptual system that is radically unreliable with regards to certain kinds of contents. For instance, suppose that the subject is a prey animal of a particularly skittish nature (for instance a rabbit). This rabbit has a perceptual "danger detector" that produces a vast amount of false positives.[99] I find it quite plausible this rabbit, if it forms beliefs of the kind "there is something dangerous over there," is not warranted in these beliefs due to the unreliability of its perceptual system. However, we can ignore this intuition for present purposes. Suppose that the rabbit is (without its knowledge or having any evidence of the procedure) captured and transferred to a far more dangerous environment that is teeming with predators. In this new environment, a high ratio of the shadows and moving things are indeed dangerous and to be avoided so its "danger detector" is highly reliable. It seems that, if the rabbit was not previously warranted in its danger beliefs, it will not *become* warranted simply because it has been transported to a significantly less safe environment. I think it unlikely that the rabbit gains the ability to form warranted beliefs by being moved to an environment in which its perceptual systems are reliable.

[98] Unless, of course, she discovers inductive grounds for the beliefs in question by noting that when she believes something to be true because she wants it to be true that belief is invariably true.

[99] My source for this case is Burge (2003).

Though Norman, Truetemp, and our skittish rabbit all end up with beliefs that are reliably true, their belief-forming processes are reliable for the wrong reasons. It is because of this feature of the cases, I suggest, one might conclude that their beliefs are unwarranted. The reliability of a belief-forming process can be accidental, lucky, or contingent in a way that undermines its ability to warrant beliefs. On the other hand, Nyrmoon, and Vipertemp seem capable of forming warranted beliefs. This is because, in some sense yet to be spelled out their belief-forming processes are reliable for the right reasons. The intuitive challenge posed by our differing intuitions about these cases is to identify the additional conditions a reliable system must satisfy in order to produce warranted beliefs. Call this general class of worries **the Contingency Problem**.

I will propose that both the Contingency and Domain Problems can be addressed in the perceptual case by recognizing that the competence of a subject's perceptual system, (and what this competence entails about the relation of the subject's perceptual system to the contingent features of a particular environment), is relevant to our epistemic evaluation of its basic perceptual beliefs. But before developing this view, we must first examine the problems perceptual systems, in particular the visual system[100], must solve in order to produce veridical representations of the distal environment and how the visual system manages to solve these problems.

[100] In the present chapter, I focus entirely on the special case of vision. Much of what I say here with regards to perceptual competence can be applied to other sensory modalities such as hearing, smell, and touch. However, it is important to note that the details of the epistemology of these other senses will depend heavily on how these senses operate and the degree to which the sub-personal processing associated with these senses resembles that of the visual system.

3. The Underdetermination Problem

In order to be able to represent its environment in a reliably veridical manner, the human visual system has to solve a host of difficult problems. Simplifying a bit, the visual system must produce a representation on the basis of the stimulation of the cells in the retina.[101] However, these proximal retinal stimulations radically underdetermine the state of the distal environment. The central question that vision science has sought to answer in the last 150 years is how the visual system is able to produce reliably veridical representations on the basis of the seemingly impoverished retinal information. This is what is known as *the underdetermination problem* or *the inverse problem* (Palmer 1999, 20). In order to explicate the notion of perceptual competence that I will rely on in my explanation of perceptual warrant, it is necessary to understand this problem, and how the visual system is able to solve it.

One can begin to see the scope of the problem by noting a provable result from a projective geometry. This is a domain of mathematics that studies how higher-dimensional spaces can be projected or mapped onto lower dimensions; for instance how information from a three-dimensional space can be mapped onto a two-dimensional plane. This makes this sort of mathematics ideal for understanding the way in which the layout of the three-dimensional environment an organism inhabits could be projected onto the (curved) two-dimensional surface of its retinas. It is a provable result that if one wishes to map a three-dimensional space onto a two-dimensional plane there is only one way in which to do so. However, there are an infinite number of ways in which to

[101] The retinal state is not the only thing relevant to visual processing. The visual system also takes into account the motion of the eye, and the tension of the muscles that control it in order to produce representations of the environment. I simplify here for the sake of exposition.

expand the information on a two-dimensional plane into a mapping in three-dimensional space. This latter problem is the problem that the visual system must solve. It must take the information registered on the retina and construct a representation of a three-dimensional environment on its basis. It must produce a single representation of a 3-D distal environment from information that is consistent with an infinite number of possible 3-D configurations (Palmer 1999, 20-21).

The problem of geometrical projection is just one of the difficulties that the visual system must deal with in order to produce a veridical representation. Another set of problems arise because the visual stimuli ambiguous in other ways. These two photos below present one such case.

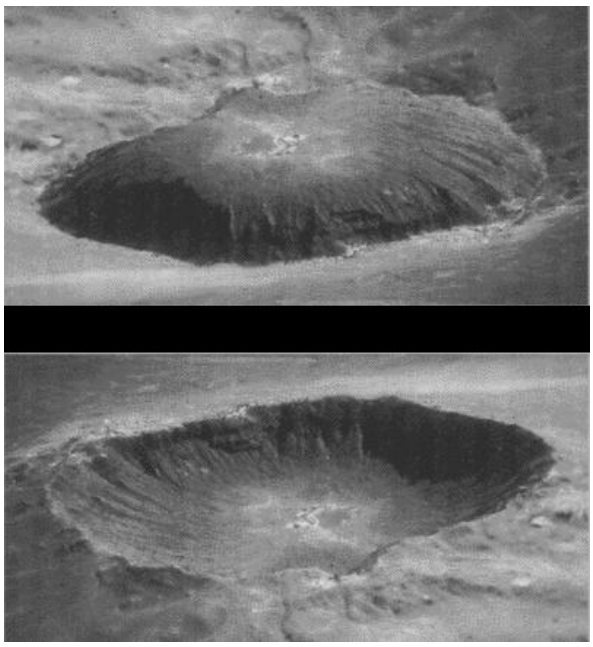

Call the top photo A and the bottom photo B.[102] In A, you will see a mound or a hill, while in B you will see a crater. However, B is merely A inverted. To convince yourself of this, simply flip the page upside down. The two photgoraphs will seem to shift so that what previously appeared to be a crater now appears to be a hill, and vice-versa. What this shows is that prior to flipping the page over, the stimulation of your retinas caused by both A and B was consistent with both the possibility that the image was of a crater, and that it was an image of a hill. Your visual system must somehow decide between these possibilities in producing a representation.

In the Hollow Face Illusion depicted below, the center block appears to have two convex faces.[103] However, the face on the right is actually a concave version of the face on the left, as the two identical blocks on either side demonstrate. The retinal stimulation you receive from the center block is consistent with both the convexity and concavity of the right face. Your visual system "chooses" between these possibilities in producing a representation of just one of them.

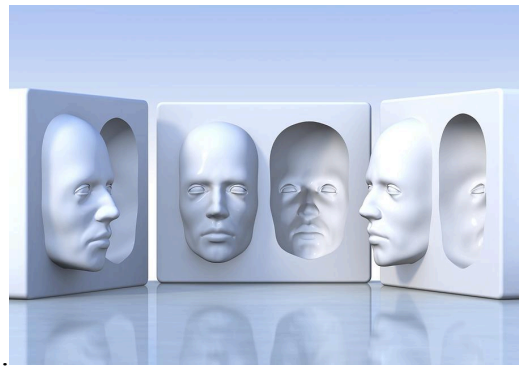

[102] Image from the homepage of Amos Storkey (Storkey 2015).

[103] This image is a photograph taken by David Mack (2013).

Your visual system must also be able to represent the state of the distal environment as remaining the same even when the proximal stimulation of the retina varies. For example: the color of an object may remain the same even when it is illuminated or shaded inconsistently across its surface; objects typically are not seen as growing in size as they approach, though they do project a larger image on one's retina as they get closer; and coins do not appear to change shape when one views them from a different angle, but the resulting retinal image will be elliptical when viewed from the side, and circular when viewed from directly above (Palmer 1999).

One's visual system must be able to represent these features of distal objects as remaining the same in spite of the variation or changes in the retinal stimulation. Indeed, the human visual system is quite good at producing veridical representations of this kind. Such capabilities are called **perceptual constancies** (Burge 2010). The images below[104] depict a well-known illusion that takes advantage of one way in which the visual system "ignores" changes in the relation of a subject to the seen object to produce stable representations of its size and shape.

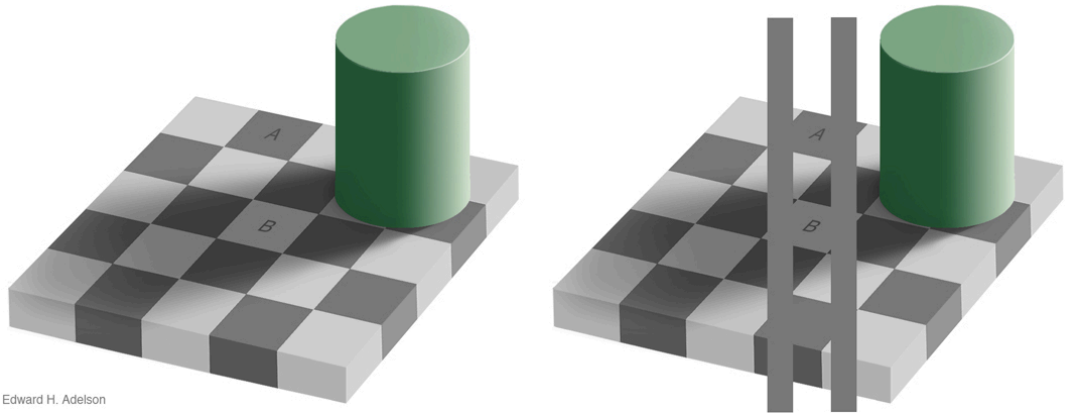

[104] This image was created by Edward Adelson (1995).

The squares labeled “A” and “B” in the left image are the same shade of grey (as demonstrated by the image on the right). However, your perceptual system represents B as a light grey square in a shadow, while it represents A as a darker grey non-shadowed square. I will return to the elements of perceptual processing this image exploits presently.

All of these images are cases of *perceptual illusions* in which the visual system fails to arrive at a veridical representation based on the information presented to it. These cases serve to illustrate, with immediate phenomenal force, the fact that the visual system must select from a range of possible representations consistent with the proximal retinal stimulation it receives. What is astonishing, when one realizes the depths of this problem, is just how successful our visual systems are at this incredibly difficult task. While we are sometimes fooled by illusions of various kinds, a normally functioning human visual system successfully represent the distal environment correctly far more frequently than it gets it wrong. We are typically correct about the color of a surface, even when the surface is lit and shadowed in complex and confusing ways. We are usually approximately correct about the relative size and distance of objects, and about their shapes even when we view them from a variety of angles. We normally have no problem distinguishing holes from hills, or convex from concave objects. Despite the radical underdetermination of the information available from the proximal stimulation of the retina, our visual systems succeed in representing our environment in a reliably veridical fashion. This underdetermination problem is the most fundamental problem in vision science: how does the human perceptual system succeed in reliably representing the distal environment given that the information it uses to do so is consistent with an infinite number of possible distal states?

4. Solving the Underdetermination Problem

i. Helmholtz's Insight

An adequate solution to the underdetermination problem should explain both how our visual systems manage the feat of reliably representing distal affairs veridically as well as why they fail in cases of illusions such as those above. In Volume III of his masterwork, *Psychological Optics* (1897), Helmholtz provided an answer that (with some revisions to be discussed below) is still the most widely accepted theoretic framework for the study of visual perception.[105] Helmholtz argued that the visual system solves the underdetermination problem by making *implicit assumptions* about the usual conditions of the distal environment and by engaging in *unconscious inferences* based on those assumptions and the proximal stimulus to produce reliably veridical representations of the environment (Gregory 1966, 2). These assumptions typically have to do with regularities of the distal environment.

In the environment in which our visual systems evolved, for instance, the primary source of illumination is the sun. Because of this, the visual system typically assumes that the light illuminating the distal environment comes from above. This allows it to eliminate a host of distal states consistent with the retinal stimulation in which the primary source of illumination comes from different directions. The Crater Illusion above is a case in which, by making this assumption, one's visual system fails to deliver a reliably veridical

[105] For more classical instances of a Helmholtzian approach see Gregory (1966) and Rock (1985). Marr (1982) is a famous version of computational theory of vision in the Helmholtzian tradition. Finally, as we will see in the next section, Bayesian models of perception such as those advocated by Feldman (forthcoming), Knill, Kersten and Yuille (1996) Maloney and Mamassian (2009), Mamassian et. al. (2002), and Rescorla (2013) are contemporary successors to the general framework that Helmholtz developed.

representation. In the top image, the perceptual system assumes that the environment is being illuminated in a particular way. Because of this, it incorrectly represents a convex mound or hill in a right-side up image rather than a concave crater turned upside down. Despite failures of this kind, the assumption that the source of illumination comes from an overhead source tends to be true in our normal environment, and allows the visual system to reliably represent the state of the distal environment based on the incomplete information in a given pattern of the retinal stimulation.

Another set of assumptions the visual system utilizes to solve the underdetermination problem has to do with variances in the illumination and reflectance properties of the surfaces of objects. Object surfaces, such as the surface of my desk, can have significant variation in reflectance properties across different parts of their surfaces. A single surface can vary in the way it is illuminated, its reflectance, color, shading, and texture. These variations are consistent with a wide range of retinal stimuli and, furthermore, a particular retinal stimulus is consistent with a wide range of distal causes. As a result, a particular pattern of retinal stimulation is consistent *both* with the variation of the color of the surface of my desk as well as a variation in the shading or illumination of my desk. The pattern is consistent with a particular element of my visual field being an edge of the desk, a spot where the desk radically shifts in color or texture, or a place in which there is a shadow that covers that area of the desk. In order to succeed in producing a veridical representation of the desk, a visual system must be capable of evaluating which properties of the surface of my desk remain constant across its surface (e.g. its color) and which change (e.g. the way it is illuminated).

How the perceptual system accomplishes this task is not a simple matter. I will focus on an element of perceptual processing that is relevant to the Checkerboard Illusion displayed above. It is a contingent fact about environments like ours that changes in the color of a particular surface tend to be relatively sharp, while changes in the illumination or the beginning of a shadow tend to be relatively diffuse. The visual system uses this fact to distinguish between edges, color-shifts and shadows by assuming that a diffuse change in proximal stimulation is due to a change in illumination, while a sharp change is due to a change in color, or to the edge of the surface. In the Checkerboard Illusion, your visual system interprets the diffuse conical pattern extending from the cylinder as a shadow rather than a change in the color of the surface, while concluding that the sharp lines of the squares on the board constitute differences in color. Because of the sharp boundaries between the squares inside the "shadow" your visual system concludes that the difference between B and the surrounding squares is due to a difference in color rather than a difference in illumination. So it represents B as a square of a lighter shade of grey in a shadow. The edges of A are also represented as differences in color rather than illumination, but there are no cues indicating that A is shaded, so it is represented as a darker shade of grey under direct illumination. Thus, your visual system represents A and B as different shades of grey when, in fact, they are the same shade. Again, the appeal to sub-personal assumptions and inferences is able to explain both the reliably successful performance of our visual system in cases like this, as well as its failure in the case of a deliberately constructed illusion (Burge 2010, 351-354).

The Hollow Face Illusion is produced by upsetting another of the visual system's assumptions about the typical state of the environment. As it happens, our natural environment tends to have many more objects that are convex rather than concave. The default assumption of the perceptual system given an ambiguous stimulus is that the distal cause is a convex surface rather than one that is concave. This is why when presented with the two stimuli in the illusion; the system represents two convex faces that are differently illuminated rather than two concave faces, or (as is the case) one convex and one concave face (Gregory 1966, 207-208).

There are several points of clarification to make about these diagnoses of the three optical illusions discussed above. First, things are far more complicated, even in these artificial cases, than my discussion here indicates. In all of these cases, the perceptual system is using a host of cues, assumptions, and inferences to produce the visual experiences that we experience.[106] I have simplified my explanations in the interest of clarity and in bringing to the fore the features of the functioning of our perceptual system that I take to be epistemically relevant.

Second, it should be emphasized that the assumptions that I have discussed here are *defeasible*. The human visual system is, of course, perfectly capable of representing the primary source of illumination as coming from other directions rather than from overhead, changes in color that are diffuse, and concave faces in the right conditions. The implicit assumptions the visual system makes about the environment can be overridden in particular

[106] For instance, the perceptual system is also highly sensitive to *patterns* of stimulations. The Checkerboard Illusion could be explained, in part, by a more holistic evaluation of the image in which the pattern of checkered squares is recognized and used to arrive at a particular conclusion about the nature of the image.

cases due to additional information available to the system. The visual system's inferences take into account *all* of the relevant information available to it to produce the representation that is made most likely to be accurate given the proximal stimulation and the background assumptions the system makes about the distal environment. So, for instance, while the perceptual system tends to produce representations of scenes illuminated from overhead, in many cases the best explanation of a given proximal stimulus will involve a light source emanating from a different orientation.

Finally, none of these assumptions or the subsequent abductive inferences are typically accessible to the perceiving subject (Burge 2010, 94). None of the computational or inferential processes that the visual system utilizes to produce a perceptual experience are something that a subject is introspectively aware of. One could be quite shocked to for instance, one could be quite surprised to find out that one's visual system assumes that objects are more likely to be convex, or that light typically comes from overhead. These assumptions could make sense on conscious reflection, but they may not ever be actively considered by the perceiving subject. Furthermore, your perceptual experiences are not sensitive to subject-level rational revision. Perceptual illusions are typically not subject to revision or change by the background beliefs or voluntary rational inferences of the subject. Regardless of the information you receive about the illusion and how it is created, your perceptual system will still represent the distal state non-veridically.

ii. Perception as Bayesian Inference

Let us turn to a contemporary successor to the Helmholtzian approach: Bayesian modeling of perception. Bayesian decision theory provides a powerful mathematical framework for modeling decisions under conditions of uncertainty. My visual system is highly reliable in producing veridical representations of the distal environment based on proximal stimuli that radically underdetermine the distal state. Thus, the central hypothesis of Bayesian models of perception is that my visual system accomplishes this by updating from prior probabilities and likelihoods in accordance with Bayes' Rule. Bayes' Theorem (on which this rule is based) can be represented as follows:

Bayes Theorem: $p(h|e) \propto p(h)\ p(e|h)$

Suppose that h is a hypothesis and e is some evidence. Bayes' theorem states that the probability of the hypothesis given the new evidence is proportional to the prior probability of the hypothesis itself multiplied by the prior probability of the evidence being present given the truth of the hypothesis. "Probability" here should be understood as subjective values assigned to h and e by an individual or her subsystems, and not as objective features of the world (Rescorla 2013). That is, these probabilities should be understood as encoded or represented by some subject rather than as a corresponding to objective regularities in the physical world. As I will understand it here, Bayes' Rule (in distinction to the theorem stated above which, strictly speaking, does not entail any normative constraints) is a normative rule on reasoning under conditions of uncertainty. Bayes Rule maintains that

given some new evidence, one ought to replace one's prior probability in a given hypothesis with the calculated posterior probability: p(h|e) (Rescorla 2013).

Bayesian perceptual psychologists hypothesize that perceptual systems like the human visual system solve the underdetermination problem by engaging in Bayesian inferences from prior probabilities and likelihoods on the basis of proximal stimulations.[107] This provides a rigorous mathematical framework and robust empirical research program that is a spiritual successor to the approach to vision science pioneered by Helmholtz (Ibid.). Bayesian models of visual perception posit that the visual system assigns a set of prior probabilities to various distal states of the environment, *h*, and prior likelihoods to *e* and *h* pairs where the *e's* correspond to some proximal retinal stimulation. Given a particular proximal stimulation, the visual system calculates a posterior probability for the *h's* and uses these posterior probabilities to produce a determinate representation (Ibid.). Updating in accordance with Bayes' Rule causes the values of the probabilities assigned to the various possible distal states to change. However, our visual experience is usually as of a *single* distal state and does not typically seem to represent a range of possibilities and the

[107] I take a realist approach to Bayesian models of perception in what follows. That is, I assume that Bayesian models of perception posit neurological states that are accurately described as prior degrees of credence in propositions—or "assignments" of subjective probability to those propositions—and that these neural states (i.e priors) are causally instrumental in the probabilistic inferences (i.e. transitions between neural states) Bayesian models attribute to perceptual systems. I assume, that is that priors are actually "encoded by" or "realized in" the biological system Bayesian theorists are modeling and that Bayesian updating on the basis of these priors is similarly encoded by (or neurologically realized in) biological processes. For an alternative, more "instrumentalist" approach, see Orlandi (2014).

probabilities associated with them.[108] Therefore, our visual system must use some means of selecting a particular *h* as the content to be represented as actual by the resulting visual experience. This last step is usually modeled by use of a *utility function*. The utility function selects a single hypothesis from the hypothesis space on the basis of some features of the posterior probability distribution (Knill et. al. 1996).[109]

The operating assumption of these approaches is that human visual system approximates an ideal Bayesian observer (Geisler, Kersten 2002). An ideal Bayesian observer updates perfectly in accordance with Bayes' Rule on the basis of priors and likelihoods that are coherent (i.e. conform to the Kolmogorov probability axioms) and mirror those regularities of the distal environment theorists have recorded to date. Though it is unlikely that our perceptual systems are ideal in this sense, in many cases our perceptual performance approximates the Bayesian ideal in surprising ways. For instance, Weiss et. al. (2002) have shown that an ideal Bayesian observer would be subject to motion illusions

[108] In reviewing a previous draft of this chapter, Aaron Zimmerman points out that one need not appeal to utility functions in order to account for the introspectable features of our visual experiences. For instance, it could be that visual processing typically results in a very high probability for one h, and that this explains the univocal nature of our visual experiences. For the purposes of the view that I defend here, such a claim would simplify matters considerably; so I am not, in principle, opposed to a Bayesian perceptual model that does not make use of such a function. However, the Bayesian models I am acquainted with *do* incorporate utility functions (see Knill, Kersten and Yuille (1996), Mamassian et. al. (2002), and Rescorla (2013)). My statement of competent dogmatism, therefore, includes reference to such utility functions in order to remain consistent with Bayesian models of this kind. The view I present can be easily revised to accommodate the excision of such functions from Bayesian models of vision.

[109] The following provide useful summaries of the central hypotheses, assumptions, and methodology of Bayesian theories of perception: Feldman (forthcoming), Knill, Kersten and Yuille (1996) Maloney and Mamassian (2009), Mamassian et. al. (2002), and Rescorla (2013). Geisler (2008) is also useful for an overview about the study of properties of natural scenes and their relevance to the Bayesian approach to the study of vision.

similar to those experienced by humans. Many of these illusions are idiosyncratic and so may seem to be due to some intrinsic fault in the construction or operation of our perceptual system. For instance: a thin rhombus moving horizontally at low contrasts is seen to move diagonally, while at high contrasts it is seen to move as in its actual horizontal pattern. However if the rhombus is 'fat', it is seen to move horizontally regardless of the contrast (Geisler, Kersten 2002, 509). Weiss and his collaborators were able to show that an ideal Bayesian observer, encoding plausible priors about the relative likelihood of fast vs. slow objects would be subject to the same sorts of illusions. This surprising result indicates that our perceptual systems at least approximate the behavior of an ideal Bayesian perceiver with regards to these sorts of cases.

The general picture one finds in Bayesian models of visual perception is of a system making use of statistical regularities of a particular environment in order to calculate the relative probabilities of various hypotheses about that environment, and then using some decision function to select a particular hypothesis based on the posterior probabilities. In what follows, I will argue that such systems are normatively evaluable in a number of ways with regards to the aim of producing veridical representations. If a perceptual system succeeds in satisfying these norms to a sufficient degree then this manifests a representational competence. If a person (or non-human animal) bases its perceptual beliefs on the outputs of a perceptual system that is competent in this sense, these beliefs are "warranted" in the sense defined above.

5. Perceptual Competence (or How to Epistemically Evaluate a Perceptual System)

i. Competent Dogmatism

There are four ways in which a Bayesian perceptual system can be evaluated with regard to its failures or successes in producing veridical representations. First, a perceptual system can succeed, by whatever means, in reliably producing veridical representations of a particular environment. In Section 2, I argued that simply being reliable in this sense is not sufficient for immediate perceptual warrant. Two primary problems emerged from that discussion: (1) the Contingency Problem, which arises out of the fact that belief-forming methods can be reliable in unprincipled or lucky ways that intuitively undermine their ability to produce warranted beliefs; and (2) the Domain Problem, which demanded a principled way of fixing the relevant domain of evaluation for the reliability of the system. Any fallible belief-forming method will be reliable with regards to some environments and not others. A proper explication of the conditions for perceptual warrant that takes reliable veridicality to be a necessary condition ought to provide a principled way in which to identify the environment that is relevant to evaluation of the competence of a subject's perceptual system.

The second way in which we can evaluate perceptual systems focuses on the prior probabilities and likelihoods that the system encodes in order to solve the underdetermination problem. The typical adult human visual system assigns higher prior probabilities to the primary source of illumination being overhead than not and to the hypothesis that a given object is convex rather than concave. A visual system succeeds, in some way, insofar as it encodes prior probabilities and likelihoods that match or

approximate actual distal regularities of its normal environment.[110] Distal regularities like the high ratio of convex to concave objects are contingent features of the natural environment in which our visual system evolved and may not be duplicated in other environments. It is possible that plentiful plantlike organisms on some distant planet have evolved to be largely concave in shape, vastly increasing the regularity of concave shapes in that environment. A cave environment could be filled with a ground-dwelling phosphorescent fungus such that the primary source of illumination there is typically from below. Atmospheric conditions on other planets could affect the diffusion of light in such a way as to render the assumptions used to distinguish surface color and illumination changes less useful. In environments like these, the implicit assumptions of our perceptual systems are more frequently false; this could (though it need not depending on the other visual cues available) result in our visual systems being far less reliable in these environments than in our own. That our visual systems encode priors that are approximately mirror the contingent state of our environment is a non-trivial achievement with regards to the whether it reliably produces veridical representations.

The third way in which the visual system can be evaluated is with regards to the inferential procedures that it employs. According to the models of perceptual processing that I am considering, the relevant norms of evaluation are clear: a perceptual system is competent to the extent to which it coherently represents the probabilities of various possible

[110] The contemporary science associated with an examination of the properties of the visually relevant regularities of the distal environment in which we find ourselves is known as Natural Scene Statistics. Researchers in this field use advanced algorithms to survey massive numbers of high quality images of natural scenes searching for regularities. These regularities in natural scenes are what the visual system exploits in order to solve the underdetermination problem (Geisler 2008, 172).

hypotheses concerning the visible properties and relations in a given organism's relatively proximate environment and, in response to changes in retinaly stimulation, redistributes these probabilities in conformity to Bayes' Rule. A maximally competent perceptual system will always redistribute probabilities on the basis of some proximal stimulus in the manner that Bayes' Rule demands.

Finally, we can evaluate the utility function that a perceptual system utilizes to select a single *h* from the probability space to be represented in experience. A utility function can be better or worse with regards to the production of veridical perceptions. Recall that Bayesian updating of the probability distribution assigned to each hypothesis only results in a new probability distribution, there is no direct selection of a particular hypothesis for representation as actual at the level of conscious visual experience. However, given that our visual experience is usually as of a *particular* distal state of the environment, the visual system must have some means of selecting a particular hypothesis from the redistributed probability space to represent. The utility functions are also normatively evaluable with regards to their production of veridical representations. For instance a utility function could fail in this regard in many different ways, for instance: by always selecting the possibility with the third *lowest* posterior probability for representation at the level of conscious visual experience, or by selecting an *h* at random ignoring the probability distribution computed at earlier stages of visual processing entirely.

A very simple and straightforward idea is to prefer utility functions that always select the hypothesis with the highest posterior probability (Mamassian et. al. 2002, 20). However, this may fail in minimizing the possibility of non-veridical representations for certain

probability distributions. Suppose that h_1 and h_2 both receive very high posterior probability assignments with h_1 getting the slight advantage. Suppose further that all the hypotheses most similar to h_1 in content receive very low probabilities, while those most similar to h_2 have very high posterior probabilities. In such a case, h_1 is very likely to veridical, but if it is not exactly correct, it is unlikely to approximate the actual state of the distal environment. h_2, on the other hand, has the benefit of having a high probability and a certain degree of safety. That is, if h_2 gets the distal environment wrong in some way, the other elements of h_2 are still very likely to be correct. It seems clear in such a case that an effective utility function should prefer h_2 to h_1 despite the fact that h_1 has the highest posterior probability.[111] Perhaps no actual case of visual processing has these features. However, what this case illustrates is that it is not certain *a priori* that the most competent utility function will simply select the hypothesis with the highest conditional probability. Thus, the desiderata for the competence of those visual mechanisms that compute and implement a utility function may be complex. But in the present epistemic context—where the focus is on the warrant of our perceptual beliefs—we can largely ignore these complexities. Simplifying substantially, let us say that a visual content fixation mechanism (modeled by a utility function) is competent to the degree it succeeds in minimizing the possibility of visual error, while maximizing the possibility of veridical visual representation.[112]

[111] This example of a probability distribution is due to Michael Rescorla. I learned much on both Bayesian probability theory and contemporary cognitive science from two of his excellent seminars on these topics in 2012 and 2014.

[112] It is surprisingly difficult to locate discussions of the relative merits of various utility functions for the modeling of perceptual systems. Mamassian et. al. (2002) present an overview of some of the issues involved, and Brainard and Freeman (1997) provide a

A competent perceptual system (that operates according to similar principles as the visual system), as I will understand the term, succeeds in all four of these ways. I will say that a perceptual system is competent if and only if:

1. The perceptual system reliably produces veridical representations in its normal environment.
2. The perceptual system encodes prior probabilities and likelihoods that mirror distal regularities of the organism's normal environment.
3. The perceptual system updates its assignment of probabilities to hypotheses on the basis of retinal stimulations in accordance with Bayes' Rule.
4. The utility function the organism's perceptual system employs to isolate a single hypothesis maximizes veridicality and minimizes error.

The thesis that I wish to put forward and explore is this: a perceiver is immediately and defeasibly warranted her basic perceptual beliefs if she forms them on the basis of perceptual experiences that are produced by a competent perceptual system. I will call this claim **competent dogmatism**.

I have stated my central thesis in a Bayesian framework (and will continue to speak in these terms in what follows). However, it is worth noting that the central ideas behind competent dogmatism can be accepted by a wide range of views about the nature of perceptual processing. I will not argue for the Bayesian approach to perceptual modeling here;[113] my focus will be on the epistemic evaluation of systems that work in the manner that models of this kind describe. However, conditions for perceptual competence can be developed in the terminology of any broadly Helmholtzian theory of the perceptual

---

highly technical discussion of the statistical merits of several competitors. The exact details of these accounts, while interesting in their own right, are somewhat beyond the bounds of my present investigations.

[113] For such arguments see: Feldman (forthcoming), Kersten and Yuille (2003), Knill et. al. (1996), Knill and Whitman (1996), Mamassian et. al (2002), and Rescorla (2013).

psychology. In more general Helmholztian terminology, a perceptual system will be competent just in case:

i. It produces reliably veridical representations of the organism's normal environment.
ii. It makes assumptions that are ordinarily true of the organism's normal environment
iii. It employs heuristic/computational/inferential procedures that are reliably truth-preserving.

These are rough conditions that are to be tailored to one's preferred account of how a particular perceptual system solves the underdetermination problem. Depending on the particular model of perceptual processing that one endorses, the details of what is required of a competent perceptual system will vary. However, so long as one's perceptual processing is broadly inferential and encodes information about the distal environment, it will be evaluable with regards to norms much like those I have discussed here. Therefore, I the success of my general theory of perceptual warrant does not stand or fall with the explanatory power of Bayesian theories of perception. Any theory of perception that is broadly Helmholtzian will be capable of formulating a suitable definition of perceptual competence adapted to the terminology of that theory using the rough generalizations provided above.

ii. Individuating the Normal Environment: The Domain Problem Solved

In my terminology, the *normal environment* for a given subject is the environment that is the direct cause of the fixation of the values of the prior probabilities and likelihoods that her perceptual system encodes. Typically, this will be the environment in which the organism's

perceptual system evolved, the environment in which the organism itself has been raised and/or the environment in which it resides. The normal environment is relevant because it figures in an explanation of why the perceptual system fixes its priors in the way that it does. Our visual systems, for instance, took millennia to develop and did so by playing an essential role in the actions are reactions of organism's causally interacting with a particular set of distal environmental states. Through random genetic variation and varying reproductive success as the result of selective pressures, the visual systems of persisting species (or populations of organisms) came to encode "assignments" of prior probability to various hypotheses about the visible aspects of their environments that "mirrored" (or accurately represented) persistent regularities in that environment. These priors became a part of the way in which the visual systems of these organisms interpreted information from the proximal stimuli it received.[114] As I have emphasized, these environmental regularities are contingent. There is nothing necessary about the prevalence of overhead illumination sources, the high ratio of convex to concave objects, or the fact that changes in color across a surface tend to be sharp. There is an evolutionary and developmental story linking the assumptions and inferential procedures our visual system utilizes in producing reliably

[114] In the present context I remain agnostic about whether or not reliably veridical representation can be *selected for* in the sense that figures in evolutionary stories. Burge (2010) argues that natural selection does not have a mechanism to single out "reliable veridicality" as a feature by which selection occurs. He acknowledges that *in some sense* reliable perceptual representations can be beneficial for an organism, he maintains that the selection must be carried out on more basic features of the perceptual system and its functioning. Graham (2011) disagrees, arguing that it can make perfect sense to speak of reliable veridicality as contributing to fitness in the relevant biological sense. I intend my discussion of causal significance of a particular environment in shaping the functioning of our perceptual systems to be consistent with Burge's more conservative view and do not assume as a necessary part of my account that the reliable veridicality of a perceptual system must be selected for in order for the system to be competent.

veridical representation to interactions with regularities of a particular kind of environment. Without these interactions, our visual systems would not encode the priors and likelihoods that they do.

Environments with different regularities from those in which an organism's perceptual system evolved are not relevant for the evaluation of a perceptual system's representational competence because they do not play the same causal role in the development of the assumptions and inferential procedures it utilizes. This is so whether or not these other environments are ones in which the perceptual system would produce reliably veridical representations if the organism were transported to them. If the organism is unreliable in these environments, this does not undermine the competence of its perceptual system because the system did not develop through causal interactions with the features of the environment that explain its unreliability.

If a perceptual system does happen to be reliable in an abnormal environment, its reliability will be accidental. By hypothesis, this type of perceptual system did not develop by interacting with the conditions and regularities of such an environment. That a perceptual system gets things right in an environment radically different from the one in which it evolved is, no doubt, fortunate for any hapless possessor of such a system that has the misfortune of being transported there. However, an abnormal environment of this kind plays no causal or explanatory role in the development of the perceptual system, its assumptions, or its inferential methods. Therefore, the competence of its perceptual system ought to be evaluated with regards to the environment in which it developed, and with which it has causally interacted with in order to fix its inferential processes.

It is plausible that many environments are causally relevant to the evolutionary development of any particular perceptual system. For instance: some of our distant ancestors likely lived in the ocean where the environmental conditions were quite different. There may have been times in the distant past (say after a massive meteor strike or volcanic eruption) when for a prolonged period the light from the sun was highly diffuse due to a heavy layer of dust in the atmosphere changing illumination conditions severely. If one is permissive enough in individuating environments, the conditions immediately following the Big Bang could constitute an environment that is causally relevant to the fixing of my perceptual system's priors and likelihoods.[115] It is possible that at least some of these ancient environments are different enough that our perceptual systems, as they are currently constituted, would be less successful at producing reliably veridical representations (than, for instance, our ancestors' systems were, or than our own perceptual systems are in environments like ours) if placed in them. Though these environments are part of the causal chain that produced our perceptual systems, it should be clear that they do not play any role in our evaluations of the representational competence of such systems. The question of whether or not my visual system would reliably produce veridical representations in a primordial ocean or in the conditions immediately following the Big Bang should be irrelevant to our evaluation of its competence.

[115] Gibson (1979) argues that an *environment* in this context should be distinguished from a set of physical states. An environment, in his terminology, is something that an organism *perceives* and *behaves in*. I am happy to take this constraint on board, but this will not help in the present context. Our evolutionary ancestors and the environments they inhabited were causally relevant to certain features of the human perceptual system.

Therefore, I claim that it is the environment that is the *direct cause* of the fixing of the prior probabilities of a perceptual system that is relevant to its normative evaluation. The distinction between direct and indirect causes is necessary to preserve intuitive distinctions in causal structures.

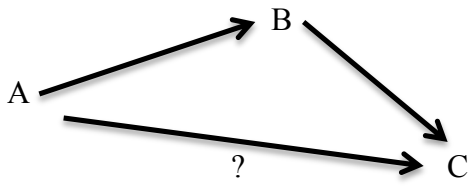

In the diagram above let A, B and C be events and let the arrows represent causal relations between the events. The question mark under the bottom arrow indicates that the causal relation here is not yet specified. Suppose that we know that A caused B and that B was sufficient to cause C. If A was the sole cause of B and B the sole cause of C then in many contexts it is perfectly appropriate to state that A caused C in virtue of its causal influence over B. However, we may still wonder if A alone sufficed to cause C. That is, we could wonder if A caused C independently of whether B happened. Suppose for instance that we held A fixed while intervening in such a way that B did not occur. Would A have still caused C in the absence of the intermediate cause B? Or did A's causal influence over C depend on the intermediate event B? If the former is the case, I will say that A is a *direct cause* of C, whereas if the latter is the case A is an *indirect cause* of C.[116]

This distinction allows us to single out the environment that is relevant to the epistemic evaluation of an organism's perceptual system from the set of environments that are causally relevant to the setting of its priors and its computational structure. Let us

[116] This distinction (as well as the diagram) is due to Woodward (2003).

illustrate this point with an artificial example. Suppose an organism evolves in a fairly stable environment, A, that is largely unchanged from the time of its earliest progenitors to a particular time t_1. Suppose further, that its perceptual system utilizes assumptions that are largely true of this original environment and updates its priors in accordance with Bayes' Rule, reliably producing veridical representations based on proximal stimuli. On the current view this will entail that organism enjoys perceptual competence sufficient for the formation of immediately warranted perceptual beliefs. However, at t_1 some sudden and massive change occurs radically changing the viewing conditions in large parts of the environment falsifying many of the assumptions that the organism's visual system makes and negatively impacting the reliability of its visual system. Call this new environment B. Now consider some later time t_n in which the organism in question or its descendants have adapted to the new distal conditions of B so that their visual systems once again reliably generate accurate representations of the visible aspects of the B-type environment in which they now live. Suppose that this return to reliability is due to the organism's perceptual faculties changing the probabilistic distribution of its prior probabilities and likelihoods (either through perceptual learning or natural selection) so that they come to mirror or track the distal regularities of B, rather than A. It is a consequence of this, that if the organism at t_n were magically spirited back to A at t_1 its visual system would be far less reliable than it is at present.

I think it clear that with regards to the aim of veridical representation, this organism's perceptual system is competent at t_n, and that the organism is capable of forming warranted perceptual beliefs. Which environment, A or B, is relevant to the evaluation of

this organism's perceptual competence at t_n? The organism is highly successful at representing the conditions in B, and this success is rooted in a principled causal connection between the distal regularities of the environment and the assumptions made by its perceptual system. This is so despite the fact that the perceptual system would systematically misrepresent its *original* environment, A. Clearly, aspects of A played an important causal role in the events that led to the existence and structure of the organism's perceptual system at t_n. Indeed, since the conditions of A persisted for most of the interval of time during which the organism's perceptual-system-type evolved, A is the environment with the most prolonged causal influence over the development of this kind of perceptual system. However, those regularities of A that are no longer represented by the perceptual system's priors no longer seem most relevant to our evaluation of the perceptual competence of those organisms endowed with the system in question. Rather the relevant features are those environmental regularities in B that are a direct cause of the assumptions that the organism's perceptual system utilizes to produce veridical representations. Thus it is B that is relevant to the evaluation of the perceptual competence of the organism.[117]

Using the distinction between direct and indirect causes, therefore, it is possible to select one environment from the set of environments causally related to the fixation of the perceptual system's priors. Furthermore, this selects the environment that seems most epistemically relevant to the evaluation of the system in question: the one to which the priors have been most directly attuned or adapted to represent veridically. Thus, while the ocean environment of my ancestors is causally relevant to the evolution of my visual system and

[117] Cases such as this distinguish my view from that held by Tyler Burge (2002). I discuss this issue and how it presents a problem for Burges' view in full detail in Section V.

plays a distant causal role in the fixing of its assumptions and heuristics, it is not the environment which is the *direct cause* of the current set of assumptions and heuristics my perceptual system actually uses to solve the underdetermination problem. Presumably the environment in which I currently find myself (or one much like it) is the direct cause of the current probability distribution utilized by my visual system, and it is this environment with regards to which the competence of my perceptual system must be evaluated.[118]

It should be clear by this point how competent dogmatism addresses the Domain Problem discussed in Section II. My visual system, for instance, has set its priors and likelihoods in virtue of interacting with a particular environment. It encodes the priors that it does to capitalize on the contingent regularities of *that* environment rather than others. Therefore, it is with regards to my normal environment that my perceptual system's competence must be evaluated. Other environments which are not directly causally relevant to the way in which the perceptual system functions, are irrelevant to the evaluation of my perceptual systems with regards to the norms of veridical representation. This proposal avoids the counterintuitive results that afflicted Goldman's various attempt to solve the Domain Problem. Unlike in Goldman's initial formulation, a subject with a perceptual

[118] This discussion raises a series of questions that a full explication of my view must address having to do with the individuation of environments and how to deal with cases in which various priors of a single subject's visual system have been fixed relative to multiple different environments. So, for instance, does an amphibian reside in one environment or two? What are the individuating conditions of an environment? If the light from above prior is fixed relative to one environment, while the convexity prior is fixed with regards to another, how are we to evaluate a whether a particular perceptual belief whose grounding experience utilizes both is warranted? These questions introduce serious complications, and perhaps some serious problems, for the account I develop here. I lay them aside in the interest of simplifying the initial presentation of competent dogmatism, and hope to address these additional complexities in future work.

system does not fail to form immediately warranted beliefs if suddenly transported to an abnormal environment. Though unreliable in the environment it currently inhabits, its perceptual system is reliable with regards to its normal environment. Therefore, its perceptual experiences are capable of grounding immediately warranted beliefs.[119]

My proposal also avoids the introduction of unnecessary chauvinism into our account of perceptual warrant. Recall Goldman's attempt to deal with the Domain Problem, according to which a perceptual system was to be evaluated with regards to *normal worlds*. For Goldman, a normal world is one that conforms to some privileged set of our beliefs about the general state of the world (e.g. there are physical objects, they move in such-and-such patterns, obey such-and-such causal laws, and so on). The problem with this proposal was that it entailed that the belief-forming methods of *all* subjects should be evaluated with regards to worlds like ours. This was implausible because subjects in environments radically different than ours, whose perceptual systems were reliable in those environments would fail to have warranted perceptual beliefs. Since their perceptual systems are unreliable with regards to *our* "normal" world, they cannot form warranted beliefs, even though their perceptual systems succeed in reliably representing their own environment. According to competent dogmatism such subjects will be warranted in their perceptual beliefs if their perceptual systems are competent with regards to *their* normal environment, not ours. Competent dogmatism, therefore, presents a principled way establishing the

[119] Of course, a subject can come to realize or gain evidence that they are in an abnormal environment, or that their perceptual system is now unreliable. Like all dogmatist views, competent dogmatism only maintains that believing the contents of one's perceptual experiences provides *defeasible warrant*. However, so long as the subject does not gain such evidence, she may continue to form warranted perceptual beliefs.

correct environment for the epistemic evaluations of the competence of a subject's perceptual system.

iii. A Measure of Charity and a Dose of Pluralism

In order to simplify my account in its initial presentation, I presented the conditions for perceptual competence in such a way that they could only be fulfilled by an ideal Bayesian perceiver. However, it is clear to all but the skeptic that less than ideal perceptual systems can be competent enough to support the formation of immediately warranted perceptual beliefs. Once one is inclined to be charitable in this fashion it becomes clear that there are likely to be many ways for a particular perceptual system to reliably solve the underdetermination problem, and thus many ways for a perceptual competence to manifest itself in particular cases.

It is unlikely that any actual perceptual system will conform ideally to the norms of perceptual competence described above. Actual perceptual systems will at best encode priors that approximate the fluctuating regularities in ordinary environments. A large forest fire may or a prolonged barrage of meteors may change the distribution of convex to concave objects in the environment without any change to an organism's prior probability distribution. Even if there are no such massive changes, it is unlikely that a naturally evolving perceptual system will manage to encode priors that *exactly* mirror the distal regularities of its normal environment. Still, despite some inaccuracies in the priors and likelihoods of its perceptual system, an organism can still be sufficiently perceptually competent to form immediately warranted perceptual beliefs. Perceptual competence

sufficient for warranted belief formation and—in the right conditions the acquisition of perceptual knowledge—should not require an ideal distribution of priors and likelihoods, but rather that some lower standard be met.

Furthermore, one should not expect, *a priori,* that a particular visual system will engage in updating procedures that perfectly accord with Bayes' Rule. Evolutionary pressures may succeed in producing perceptual systems that are highly reliable in veridically representing an environment . In practice, it seems that our perceptual systems do, in many cases, calculate posterior probabilities in a way that very closely approximates an ideal Bayesian inference.[120] However, a healthy dose of caution is warranted regarding the thesis that our visual system engages in computational procedures that exactly instantiate this statistical ideal. It is an open empirical question when, how, and why our perceptual systems fail to instantiate precise patterns of Bayesian updating. A plausible working hypothesis at this stage is that our perceptual systems are *very good* but not *ideal* Bayesian observers. If this is so, then an account of perceptual warrant that appeals to Bayesian norms must account for this fact.

Similar points apply to the utility function that a perceptual system uses to isolate a hypothesis it then "broadcasts" to conscious visual experience after updating its credences based on proximal stimulation. First, there could be many different processes or mechanisms—represented by different utility functions—some more appropriate than others for particular environments. That is, in some environments it could be the case that the peculiar probability distributions I discussed in the previous section are exceedingly rare and

[120] For instance, consider again the results of Weiss et. al. (2002).

a simple "highest posterior probability" selection method can be safely and effectively employed. Other environments may require a more complex utility function in order for the perceptual system to manifest competence as I have defined it. In addition, as in the case of the updating procedures, it seems that an organism might posses the requisite degree of perceptual competence even if its experience content fixation mechanism or process fails to implement whatever utility function would maximize veridicality and minimize error in the particular environment in which it is operating. Again we must allow that organisms can be perceptually competent without possessing ideal perceptual systems.[121]

Because of these considerations it is clear that perceptual competence sufficient for warranted belief-formation need only approach or approximate the ideal laid out in Section 4. However, this is only to acknowledge that perceptual competence, like warrant, justification, and many other epistemic goods, comes in degrees. According to the current proposal, a perceptual system is more or less competent to the degree that it approximates the prior probability distribution, inference procedures, and utility function that an ideal Bayesian observer would implement in the "normal" environment of the organism whose system it is. However, we should grant that a less than maximal degree of competence is required for perceptual warrant and perceptual knowledge. The perceptual systems of particular organisms will fall on a scale with some (probably vague) degree of competence sufficing for warranted belief-formation. Since warrant intuitively comes in degrees, this feature of the present account is just what we should expect from an account of perceptual warrant.

[121] Of course, even an ideally competent perceptual system will still be fallible, due to the underdetermination of its evidence.

A somewhat more interesting side effect of this acknowledgment is that there are likely to be many ways for a perceptual system to achieve the required degree of competence sufficient for the formation of warranted perceptual beliefs. A higher degree of competence in one area may compensate for deficiencies in others. For instance, a perceptual system may have systematic failures in how it updates its prior probabilities based on some stimulus that are compensated for by an exceedingly effective and well-tailored utility function. More inaccuracies in the prior distributions of probabilities could be allowed for in a system that employs a highly rigorous updating procedure that conforms precisely with Bayes' Rule. In short, there will be no one way for a perceptual system to manifest the required degree of competence to ground an organism's ability to form warranted perceptual beliefs. Given the diversity of perceptual systems furnished by evolution both on this planet and possibly on others, I take this pluralism to be a salutary consequence of the present view. The framework I propose here is quite general, in that it can apply to, and explain the perceptual warrant-conditions of the perceptual beliefs of a wide variety of organisms with perceptual systems that evolved in very different conditions, and that have been shaped by the environment to solve the underdetermination problem in many different ways.

6. Perceptual Competence and Perceptual Content

Tyler Burge (2003) offers a view of perceptual warrant to which the view I have developed here owes a significant intellectual debt. As my view is largely an elaboration of the approach that he advocates, special attention to his view, and my reasons for revising it are

warranted. As with competent dogmatism, Burge solves the Domain Problem by arguing that the reliable veridicality of an organism's perceptual experiences can be tied in a principled manner to interactions and adaptation to a *particular* environment. Unlike the view that I defend above, Burge takes the epistemically relevant interactions to involve the fixing of the contents of the perceptual experiences. On Burge's view it is the content-fixing interactions, rather than the fixing of the perceptual system's encoded priors and likelihoods, that are epistemically relevant. I will argue that this presents problems in cases where a changing environment can result in changes in the manner in which a perceptual system processes information. The ability of the position I have advocated above to deal with such cases is the primary motivation for preferring my view to Burge's.

As I have emphasized, the reliability of a perceptual system must be evaluated with regards to some environment or other. Burge recognizes that perceptual systems develop over the course of evolutionary time due to causal interactions with an environment with certain persistent regularities. This is typically also the environment responsible for the fixing of the representational contents of the perceptual system. Take two intrinsically physically identical organisms: Cracked and Shadow. Cracked's visual system developed in an environment where there is a high frequency of cracks, but very few shadows (perhaps a sandstone desert landscape). Shadow, on the other hand evolved in an environment that is *full* of shadow casting objects, but with very few cracks. Suppose that Cracked and Shadow receive an identical retinal stimulus. The stimulus in both cases is consistent with both the possibility that a certain bit of the distal environment is a crack, as well as the possibility

that it is a shadow. Intuitively, Shadow's visual system will represent a *shadow* while Cracked's will represent a *crack*.[122]

As representational states, perceptual experiences will be partly individuated by their contents. My perceptual experience as of a red ball is type-distinct from my experience as of a blue cube partly because of their differing contents. A constitutive condition on an experience being the very thing that it is is that it has a particular content. In Burge's terminology these contents are *anti-indivudalistic* (in more common terminology they are said to be *externalist* or *externally individuated*); they are fixed by contingent causal interactions with a particular environment. The reason that Cracked represents *crack* while Shadow represents *shadow* when presented with the same retinal stimulus is explained by the fact that the contents of their perceptual states were fixed in relation to different environments. Furthermore, Cracked's visual system is reliably veridical at crack-representation in its normal environment, while Shadow's is reliably veridical at shadow-representation in *its* home.

Burge concludes that the environment in which the contents of an organism's perceptual contents were fixed is the correct domain of epistemic evaluation for the reliability of its perceptual systems because this environment plays a greater role than any other in fixing the contents of the perceptual states and therein making them the kinds of states that they are. It is because other environments do not play such a role that they are irrelevant to the epistemic evaluation of the system (Burge 2003, 536). According to Burge's proposal the environment with regards to which a perceptual system should be

[122] This case is originally stated in Burge (1989)

evaluated for reliability and warrant is the environment in which the contents of the organism's perceptual states were fixed. If an organism's perceptual system reliably represents those contents in the environment in which they were fixed, then, its experiences are capable of providing immediate and defeasible warrant for its basic perceptual beliefs.

In many cases, Burge's proposal and competent dogmatism will render the same verdict concerning the epistemic status of an organism's perceptual beliefs. This is because in most ordinary cases the environment in which the contents of an organism's perceptual experiences were fixed *just is* the environment that is the direct cause of the fixation of the prior probabilities and likelihoods the perceptual system encodes. In both cases, this is plausibly the environment in which the organism's perceptual system evolved or that in which the organism resides and/or was raised. Indeed, there is a principled connection between these two processes. Suppose that in the case of Cracked and Shadow that their retinas are intrinsically identical, and that their visual systems are sensitive to the same sorts of features of the retinal stimulation.[123] Plausibly the reason that Cracked sees a crack in this case while Shadow does not is that Cracked's visual system puts a higher conditional probability on the presence of cracks given this stimulus than does Shadow's. It does so because of a contingent regularity in the environment in which its perceptual system evolved. This environment is also causally responsible for fixing the contents of Cracked's and Shadow's perceptual experiences. That is, Cracked's ability to perceptually represent cracks and Shadow's ability to perceptually represents shadows constitutively depend on their respective causal interaction with cracks and shadows. Therefore, both the contents and

[123] For instance, as Burge first described the case, they are both sensitive to zero-crossings across different sized computational filters of the retinal stimulus (Burge 1989)

the prior probabilities encoded by their visual systems were fixed by repeated causal interactions with their respective environments.

However, this need not always be the case. The prior probabilities and likelihoods a visual system encodes can change while the contents of its perceptual states remain fixed. This can happen when the contingent regularities in an environment change after the initial development of the perceptual system or after the initial fixing of the contents of the perceptual experiences it produces. Suppose that an organism evolves in an environment similar to ours in which the primary and most persistent source of light comes from overhead. The organism is of sufficient complexity to form beliefs about its immediate environment, and has a range of basic contents that it is capable of representing both perceptually and conceptually: e.g. crack, shadow, edge, motion, various sizes and distances of objects, and so on. These contents were fixed by interactions with the environment in which the organism's perceptual system evolved: namely one in which the light comes predominantly from above. Furthermore, due to this regularity in the distal environment, one of the assumptions its perceptual system uses in solving the underdetermination problem is that light typically comes from above. In this environment, the organism's perceptual system succeeds in reliably representing a range of distal states, much as ours do. Both competent dogmatism and Burge's proposal entail, quite plausibly, that such an organism can form immediately warranted beliefs on the basis of its perceptual experiences.

Suppose that due to some calamity the sun on this organism's planet is obscured by thick clouds so that virtually no illumination from that source reaches the surface and simultaneously a phosphorescent fungus becomes prevalent. In this new environment light

typically comes from below. While at first the organism's perceptual system continues to assume that light will tend to come from overhead, over time its (or its descendants') perceptual system shifts to accommodate the change in distal regularities. That is, in order to better represent the new environment, it now gives a higher prior probability to the hypothesis that the primary illumination source is below rather than above. This change in priors allows the organism to represent distal states in the post-cataclysm environment in a reliably veridical manner.

Cases like this present a problem for Burge's account. The environment in which the contents of the organism's perceptual states were fixed is plausibly the pre-calamity environment. Yet, due to the changes in the environment and the organism's adaptation to the changes, the organism's perceptual system may turn out to be less competent to represent such an environment. Its visual system is now attuned to an environment with different distal regularities, and those changes could cause it to produce non-veridical representations at a higher rate in environments like the one in which it initially evolved. While producing reliably veridical representations in the post-calamity environment, the organism's visual system would, if placed in the environment in which the contents of its perceptual experiences were fixed, be significantly less reliable. Burge's view is seemingly committed to a counter-intuitive verdict on this case: that when the organism's visual system successfully adapts to its new environment this mitigates the "warranting power" of its visual experiences. However, we are inclined to think of this change as a remarkable adaptive *success* with regards to the organism's visual system's constitutive aim of producing reliable representations. The organism's visual system has adjusted to represent

the new environment in which that organism finds itself. It is this latter environment by which we ought to be evaluating the representational successes and failures of its perceptual system, rather than the one in which the contents of its experiences were initially determined.

The case above is a highly artificial example. However, the phenomena of shifting perceptual priors has been documented empirically. To give one such example, Adams, Greg and Ernst (2004) demonstrated that one could alter the light-from-above prior by training subjects with a combination of visual and haptic stimuli. They were able to show that the subject's judgments (and experiences?) of convex images shifted to incorporate the assumption that the source illuminating those images was skewed slightly to the side. In this case, of course, it is open to Burge to maintain that the shifts in illumination direction documented by the trial were not severe, and would likely not contribute in significantly negative ways to the overall reliability of the visual system in question. However, the science fiction case presented above illustrates that such changes need not, in principle, be slight, and also suggests that in some extreme cases changing the priors probabilities a visual system encodes could result in representational differences significant enough to affect reliability in a particular environment. If this is so, then the example presents a serious problem for Burge's position.

It was reflecting on this problem for Burge's account that led to the development of the view that I have presented here. Suppose, as is likely, that the organism's pre-calamity priors and likelihoods were directly caused by the environment in which the contents of its perceptual experiences were fixed. Competent dogmatism is in agreement with Burge that

prior to the calamity the organism enjoys immediate perceptual warrant for beliefs formed on the basis of these experiences. Post-calamity, however, the priors that the organism's visual system encodes have changed. Due to repeated interaction with what would have previously been anomalous distal states, the priors regarding the likely direction of illumination have shifted to prefer the hypothesis that light comes from below rather than from above. This shift in the encoded priors of the perceptual system is directly caused by interactions with the new environmental conditions. Therefore, according to my proposal, the organism remains warranted in its perceptual beliefs despite the shift. This presents a substantial improvement to Burge's account.

7. The Contingency Problem Revisited

In this section and the next I present some implications of competent dogmatism by discussing what it says about a series of philosophical thought experiments. The present section will focus specifically on those thought experiments brought to bear in stating the Contingency Problem in Section 2. There, I described the Contingency Problem as establishing that reliability does not suffice for perceptual competence or immediate perceptual warrant. In such cases we are inclined to say that, though reliable, the subject's belief-forming method is reliable *for the wrong reasons*. In this section, I argue that competent dogmatism has the resources to provide explanations for the various intuitive epistemic judgments these cases elicit.[124]

[124] I make the distinction I do purely for heuristic purposes. All the cases I discuss bear some similarities to one another, so I don't maintain that they divide neatly into exhaustive

In some cases, an organism's perceptual system can be reliable due to some accidental or lucky features of its situation. First, consider a Benevolent Demon-Worlder (BDW) who is possessed of a perceptual system that is disposed to reliably misrepresent its normal environment. That is, given a particular retinal stimulation, BDW's visual system will be highly likely to produce a representation that is very different from the distal states that caused the stimulation. However, BDW is fortunate in that an omnipotent and benevolent demon has taken pity on her. The demon knows ahead of time how BDW's visual system will represent the world to be, and uses its power to remake the world so that it is the way that BDW perceives it to be. BDW's perceptual system is reliable, but her perceptual beliefs are intuitively not warranted.

Second, consider again our skittish rabbit and its perceptually grounded "danger-detector." In its natural environment, the rabbit has few predators and an overactive danger-detector. The rabbit perceives danger in response to shadows, any fast movement, and a host of other stimuli. Since, in the vast majority of these cases, there is nothing dangerous present the rabbit's danger-detection is highly unreliable. Suppose that the rabbit is captured and moved to a small enclosure full of wolves. Since nearly everything in this new environment is dangerous, the rabbit reliably represents danger for the remainder of its (probably short) life. It was pointed out that whatever one is inclined to say about the rabbit and its beliefs about the presence of danger prior to being moved, one should say the same thing *after* the move. In particular, even if one takes reliability to be a necessary condition

---

and exclusive categories. However, drawing the somewhat artificial line that I do helps to bring to the fore the two general features of competent dogmatism that will be appealed to in addressing problems of this kind generally.

for perceptual warrant, one should conclude that simply moving the rabbit to an environment in which it reliably gets things right does not allow it to form warranted perceptual beliefs.

It is consistent with competent dogmatism to deny that the subjects in these two cases are warranted in their perceptual beliefs. Moreover, the view suggests a natural diagnosis of their epistemic deficiencies. BDW and the skittish rabbit fail to form warranted perceptual beliefs because the relevant perceptual capacities are not competent. The danger-detector of the rabbit is unreliable in its normal environment. It may be evolutionarily beneficial to this organism to have a perceptual system of this kind. For instance, perhaps the predators in its normal environment, though rare, are quite effective and that the rabbit is sure to be eaten if it encounters one unless the rabit has a large head start. This could provide an evolutionary explanation for why the rabbit has developed an incompetent perceptual faculty. Nevertheless, with regards to the norms of veridical representation, its danger-detector still fails to be competent in its normal environment. As I maintained in Sections 2 and 5, one ought not always evaluate the competence of a perceptual system relative to its performance in the environment in which the organism currently finds itself. Rather, one must consider the performance of the perceptual system in the environment that directly causally explains its nature and function. In the case of the rabbit, this is the relatively predator-free natural environment rather than the wolf pen. In its normal environment, its perceptual system is unreliable, and so incompetent.

Similar remarks apply to BDW. Though reliable in her normal environment, this reliability is not due to any competence of her perceptual system, but rather to the heavy-

handed intervention of an omnipotent benefactor. Her perceptual system is disposed to systematically misrepresent BDW's environment on the basis of how her retinas are stimulated. This could occur for a range of reasons: perhaps the priors or likelihoods that her perceptual system makes use of do not mirror distal regularities; perhaps, her perceptual system does not derive posterior probabilities in a manner that approximates Bayes' Rule; perhaps the utility function selecting a particular *h* from the set of hypotheses is very poor; or perhaps some combination of all of these failings explains her incompetence. Whatever the reason, without the intervention of the demon, her perceptual system would not represent her normal environment reliably. In other words, the reliability of her perceptual experiences do not entail that her perceptual system is competent to represent her normal environment. Her reliability is not attributable to her as a cognitive subject. I suggest that this is the reason that we do not judge her to be warranted in her perceptual beliefs despite her reliability.

Competent dogmatism requires more than mere reliability. It requires that a subject's perceptual system manifest normatively laudable features with regards to the aim of producing veridical representations. These norms are not dissimilar to those that apply to subject-level reasoning and belief formation; the background assumptions the perceptual system uses must be well grounded, and the inferences the perceptual system utilizes must be roughly in accord with good inferential norms. The subject need not be aware that her perceptual system succeeds in these ways, but having a perceptual system with such good-making features makes a subject epistemically better than one who does not. Cases such as BDW and the skittish rabbit illustrate that we are not inclined to judge a subject to be

warranted when the reliability or veridicality of her perceptual system is due to factors that do not depend on any epistemic competence that she possesses. Competent dogmatism respects the intuitive force of these considerations, not by imposing additional, over-intellectualized subject-level constraints on warranted belief formations, (as suggested by internalists like BonJour (1985) and Lehrer (1992)), but rather by imposing normatively motivated constraints on the functioning of a subject's perceptual system.

A slightly different but related point is illustrated by considering Lehrer's (1992) Truetemp case. In this case Truetemp reliably forms true beliefs about the temperature outside of his house due to an artificially implanted device. Strictly speaking, this case cannot serve as a counterexample to competent dogmatism because this view only makes claims about the warrant for subjects' perceptual beliefs. By stipulation, Truetemp does not form his temperature beliefs on the basis of any perception. Nonetheless, competent dogmatism has the resources to explain the difference between the original Truetemp case and Lyons' (2009) Vipertemp variation. Vipertemp, recall, is a member of an alien species with a perceptual system resembling that of the pit vipers of Earth. By pointing his heat-sensing pits in the right direction, he is able to ascertain the temperature of the air outside of his dwelling. Unlike in Lehrer's original case, it seems perfectly intuitive to assert that Vipertemp is warranted in his beliefs about the temperature outside of his dwelling.

Our differing judgments in these cases can be explained, on the present approach, by appealing to fact that while the normatively laudable features of a subject's perceptual system play an important role in our evaluation of her epistemic state, the accuracy or reliability of some artificially implanted device does not. Intuitively, this is because the

competence of a subject's perceptual system is *her* competence: it is attributable to her as a subject. On the other hand, the functioning of some artificially implanted device like Truetemp's temperature detector is not. What such cases show, then, is not that externalist conditions on warrant are inadequate, but rather that these conditions must involve epistemic competencies of the subject herself rather than some device. Competent dogmatism maintains that the manner in which a subject's perceptual system produces experiences on the basis of proximal stimulations is a way in which a subject can be properly epistemically evaluated.[125]

BonJour-style clairvoyance cases have very little force here. Such cases typically trade on the "spookiness" of clairvoyance. The claim that such subjects are not warranted in their beliefs is most plausible in cases in which the clairvoyant faculty is under-described or described in such a way that there is no physical causal mechanism connecting Norman to the states of the world that make his beliefs true. In such cases, Norman's faculty seems to be no more than a series of lucky guesses that indicate absolutely no epistemic competence or principled connection to the truth. However, if one states the case along the lines of

[125] There is an interesting range of cases to consider here. Obviously *some* artificial enhancements are consistent with a subject having immediately warranted perceptual beliefs. For instance, if I unknowingly undergo laser surgery to correct my nearsightedness, I could still form warranted beliefs afterwards. In more extreme cases, one can imagine a future in which cybernetic eyes, optic nerves, and/or other sensory organs are available can be implanted to greatly improve one's perceptual capacities. In such cases, (even if the subject is unaware of the procedure) I would tend to think that she is warranted in her post-procedure perceptual beliefs. Discussing the full range of such cases here would take us too far afield, (though I find thinking about such cases stimulating). Many of these cases do not present *prima facie* problems for competent dogmatism. Corrective surgery, or even the implanting of cybernetic eyes relies on pre-existing perceptual competencies of the subject. The information is still processed utilizing prior probabilities and computational heuristics of the subject's cognitive system. All that has changed in such cases is how the perceptual system is proximally stimulated.

Lyons' (2009) Nyrmoon case, in which the reliable belief-formation is due to a competent perceptual system (admittedly of a very different kind than our own) it would amount to mere chauvinism to deny that such a subject could be warranted in its perceptual beliefs. That the organism is perceptually sensitive to different features than light or sound waves should not preclude the possibility that it can form warranted perceptual beliefs.

I do not pretend to have considered all the potential objections along the lines of the contingency-based worries. However, I hope that it is clear that the view that I am defending has substantial resources at its disposal to address these sorts of problems. This ability arises from the view's insistence that the reliability of a perceptual system be tied in a principled manner to the regularities of a particular environment. That a perceptual system encodes priors that approximate distal regularities and updates in accordance with Bayes' Rule are non-trivial accomplishments with regards to the aim of producing veridical representations. A perceptual system with these features reliably represents its normal environment and does so in a manner that manifests a competence with regards to the aim of veridical representation. That is, its success is not due to some lucky or highly contingent external circumstances, but rather to the normatively good features of its internal processing and to its development in direct causal interaction with a particular environment.

8. The Domain Problem Again: Subjects in Skeptical Scenarios

A philosopher of a certain (internalist) persuasion will think that I have failed to tackle the most significant problem afflicting *any* externalist account of warrant, perceptual or

otherwise.[126] As discussed in Section 2, these problems arise when considering the epistemic status of subjects in skeptical scenarios. In this section, I argue that the most serious issues concerning asymmetries in epistemic status between normal subjects and their counterparts in skeptical scenarios can be resolved by competent dogmatism. While my position does entail that there are certain person/BIV pairs that are evaluated asymmetrically, it is far from clear what our intuitions about such cases are or should be, and as such they do not present persuasive counterexamples to the view that I advocate.

A persistent complaint leveled against externalist epistemological theories is their seemingly counterintuitive results with regards to subjects in skeptical scenarios (Cohen 1984, Huemer 2006, Smithies 2014). Such cases are often appealed to in motivating the Domain Problem for reliability-based epistemological views. This chapter has been long, so allow me to repeat the internalist argument given in Section 2. A subject in a skeptical scenario, such as a brain in a vat (BIV), will, by stipulation, employ belief-forming methods that are highly unreliable. Thus, a reliabilist view such as Goldman's (1979) is committed to the claim that subjects in skeptical scenarios are unwarranted in their beliefs. However, embodied subjects, such as myself, (I presume), employ methods that are reliable and can, therefore, form warranted beliefs. Consider my BIV counterpart, BIV-Tim.[127] BIV-Tim has

[126] For instance: this seems to be the motivation behind Huemer's (2006) general arguments against any externalist conditions for warranted.

[127] In this initial explication of the argument, I will be rather vague about the precise features of the case. The present discussion is meant to introduce the general problem for externalist views posed by such cases. So, for instance, I will remain neutral in this initial presentation about what is required in a description of a BIV case in order to make it plausible that the two subjects share the same contents for their perceptual experiences and their beliefs. I will discuss many varieties of BIV cases in what follows, and in that context

a sequence of experiences that are introspectively indistinguishable from my own. Despite his unreliability, it seems callous to contend that BIV-Tim is unwarranted while I am perfectly warranted in all my beliefs. After all, BIV-Tim reasons just as I do, forms his beliefs on the basis of the same experiences, and is in all ways (it seems) epistemically like me. Given that BIV-Tim forms all of his beliefs for precisely the same reasons, using the same methods, and based on the same experiences as I do; if I am warranted in my perceptual beliefs, so too must BIV-Tim be.

The problematic nature of the asymmetry can be further demonstrated by considering a small addition to the case. Suppose that, while we both slept one night, BIV-Tim's brain were surgically implanted in my body and mine were put into his vat. We awake and continue to form our beliefs in the same manner that we had previously. If reliabilism is correct, then, through no fault of my own, I have lost the ability to form new warranted perceptual beliefs due to the fact that my perceptual systems are unreliable. Meanwhile BIV-Tim, through no revisions to his belief-forming practices, has gained the ability to form warranted perceptual beliefs because *his* perceptual systems are now reliable. But, given that neither of us has done anything to change our respective belief-forming procedures, it seems that there is no way in which this miraculous epistemic shift has occurred. The mere changing of our circumstances, something of which neither of us is aware, does not seem to support so radical a change in our respective abilities to form warranted beliefs. I grant that in cases like these an adequate theory of perceptual warrant out to be able to account for the epistemic similarity between such subjects. I believe that competent dogmatism is capable

---

the details of the cases will play a significant role in how competent dogmatism addresses them.

of handling cases of this kind in a way that ought to satisfy all but the most dogmatic internalists. In the most compelling versions of such cases, competent dogmatism succeeds in rendering the intuitive result that both subjects are equally warranted in their perceptual beliefs.

On the most natural reading of the case, competent dogmatism entails that BIV-Tim is warranted in his perceptual beliefs. If BIV-Tim's brain is a *human* brain, then it is a result of the same evolutionary process as mine. The operation of BIV-Tim's perceptual system has been determined by the same evolutionary pressures as have mine. Like mine, BIV-Tim's perceptual system will tend to assume that light comes from above, that distal objects are convex, and so on.[128] Moreover, given that his is a human brain the causal/evolutionary explanation of the functioning of my perceptual system applies to him as well; the environment in which his perceptual systems developed and in which the assumptions and inferential procedures they utilize were fixed is a normal physical world, indeed *our* world. As such, the competence of BIV-Tim's perceptual system must be evaluated with regards to *our* world rather than the deceptive vat world in which he finds himself. BIV-Tim's visual system encodes the same priors about the environment, utilizes the same inferential procedures, and, if placed in this environment, would be equally successful in producing veridical representations of the distal environment (as the construction of the case suggests). Since BIV-Tim's perceptual competence is equivalent to my own, it follows that, like me, he

[128] As I mentioned above, it is possible for the priors encoded by an organism's perceptual system to change over the course of its life as a result of certain sorts of stimuli. However, by stipulation my BIV duplicate has had all the same stimuli to the visual centers of his brain as I have. Thus, there will have been no opportunity for the priors and likelihoods our systems encode to diverge.

is capable of forming immediately and defeasibly warranted beliefs on the basis of his perceptual experiences. Even though competent dogmatism claims that reliability is a necessary condition for perceptual warrant, it does not entail any epistemic asymmetry between subjects and their BIV-counterparts.

Competent dogmatism does not entail that when BIV-Tim's and I are swapped that I lose the ability to form warranted beliefs while he gains it. Whether or not our basic perceptual beliefs are immediately warranted depends on whether or not our respective perceptual systems are competent. This competence is evaluated with regards to the environment in which the assumptions and heuristic procedures of the perceptual system were fixed. My perceptual competence is evaluated relative to my normal environment, not the peculiar BIV environment and this competence persists even if I have been unwittingly transferred to a radically deceptive environment. So long as I do not receive any evidence that this has occurred, I will continue to form warranted (but false) perceptual beliefs.

It will be objected that the skeptic has the entire range of metaphysical possibility at her disposal for constructing her scenarios. Surely it will be possible to construct a case in which there is a subject, who is my introspective duplicate while lacking perceptual competence. I grant at the outset that this is so. However, I hope to show that it is not at all obvious that this consequence is problematic when one considers exactly which kinds of organisms meet these criteria. I will not pretend to discuss all possible variations on the BIV case that could be relevant here. Instead, I will look at a small sampling of such cases, and discuss how the view I defend addresses each. It is my hope that the sampling that I have chosen will provide a template for how competent dogmatism would address the many

varieties of skeptical scenarios that I do not discuss. I will be discussing many such cases, some of which may be though tot be impossible based on various theoretic commitments that one might hold. Since these cases are meant to provide problems for my position, if any or all of them turn out to be impossible, this is beneficial to the position that I defend. However, I take it that at least *some* will think of each case that it is possible, and so my discussion of each is addressed to those individuals who think the case is possible and provides a *prima facie* problem for competent dogmatism.

Let us first look at what I will call *spontaneously generated cognitive systems*. These are possible cognitive systems that are generated instantaneously and/or accidentally, rather than through an evolutionary process. This can either occur due to some cosmic or quantum accident, or by the conscious actions of some agent (e.g. the development of an artificial intelligence in a laboratory). One such case is Putnam's (1981) version of the BIV thought experiment. A Putnamian BIV is a brain-like object that comes into being by an accident of quantum indeterminacy attached to a similarly created supercomputer. The computer feeds the BIV stimuli intrinsically identical to those that my own brain receives from the stimulus of my perceptual organs. Such an object is not a *human* brain, but it is an intrinsic physical duplicate of one.

According to the present proposal, if we are to evaluate such a BIV for perceptual competence, we must first identify its normal environment. In the case of a Putnamian BIV, if *any* environment is to play this role,[129] it must be the vat-environment in which it finds

[129] Davidson's (1987) case of Swampman in which a bolt of lightning striking a swamp by a quantum accident produces a physical duplicate of me is plausibly a case in which there is (at least initially) no environment plays the content fixing role. If this is correct then,

itself. As Putnam (1981) convincingly argues, given that such a BIV has no causal connection whatsoever to a non-vat environment like ours, it is unlikely that the contents of its beliefs or perceptual states refer to physical objects like *trees*, *tables*, *humans* and so on. Rather, the contents of such a BIV's thoughts and perceptual states have been fixed by interacting with (what it calls) objects in its simulated experiences. That is, the content of its perceptual experience is not *edge* but *vat-edge*, not *shadow* but *vat-shadow* and so on. The simulation is also the environment with regards to which the priors and likelihoods used by the BIV's perceptual system have been fixed. Therefore, the normal environment for such a BIV is not *our* world (with which it has never interacted) but rather the simulation world that the computer is generating. With regards to this vat-world, the BIV meets my conditions for perceptual competence: its perceptual system reliably produces veridical representations of the environment, and it does so on the basis of assumptions that are reliably true of its normal environment. Thus, my account agrees with Putnam's diagnosis of the case: a Putnamian BIV is warranted in its perceptual beliefs. Indeed, this is plausibly a case where the BIV gains *perceptual knowledge* because the contents of its perceptual experiences and its beliefs are *veridical* and entirely appropriately caused by the

---

despite Swampman's physical similarity to me, he does not possess contentful thoughts or perceptual experiences. Of course, over time Swampman could acquire these contents by interacting with an environment like mine. Swampman is another example of a spontaneously generated cognitive system. However, because of the problems introduced by his lack of perceptual contents, I do not consider the case in depth here.

environment it inhabits.[130] A spontaneously generated cognitive system of this kind presents no problem for competent dogmatism.

Let us now abandon intrinsic physical duplicates and consider the brain and/or nervous system of some non-human alien. Suppose that on some distant planet alien scientists have kidnapped one of their own infants, and placed its brain in a vat. Throughout the course of its life they have given this subject experiences as of being a human in a world much like ours: indeed by sheer chance, they have given this unfortunate alien a series of experiences that are introspectively indistinguishable from *my* experiences. That is, this alien BIV is an introspectively indistinguishable counterpart of mine. The alien BIV is not a Putnamian BIV. We can suppose that due to the evolutionary development of its species' perceptual systems its perceptual experiences have contents involving normal physical objects and not mere simulations. The alien BIV is systematically deceived, and is radically unreliable in the formation of its beliefs. But, one might be inclined to think, the alien BIV is perfectly warranted in its perceptual beliefs.

First, it is unclear whether every possible brain/nervous system is capable of feeling sensations introspectively indistinguishable from my own. It may be impossible to stimulate an alien being with a brain and ordinary perceptual systems sufficiently different from my own into having experiences like mine. For instance, an organism with no native visual system may not have the 'wiring' necessary in order to experience visual sensations so that

[130] For useful discussions of these issues see Brueckner (2013) sections 7-11. I agree with Brueckner's conclusions that while Putnam's diagnosis of the case he presents is ultimately persuasive, his general argument against the skeptic fails. Every BIV is not a Putnamian BIV, and powerful skeptical arguments can be stated in terms of e.g. recently envatted brains, or brains that have been forced to spend every other weekend in a vat while spending the rest of the time in a normal human body.

no possible stimulation of the brain would produce them. It also may be the case that the native world of this organism must bear certain similarities to a world like ours. For instance, it may be necessary for this organism's ordinary environment to include edges, shadows, illumination conditions, and shapes much like those that we encounter in our normal enviornment. Suppose, then, that our alien brain is sufficiently like our own that stimulations of its brain can produce experiences like our own. Suppose that it has perceptual systems that engage in similar sorts of processing to our own and encode prior probabilities and likelihoods about its own native environment, that these priors mirror environmental regularities of its ordinary environment, and the inferential processes its perceptual system uses roughly conform to Bayes' rule. That is, suppose that the alien brain is in possession of a competent perceptual system, though the competence is tied to a different environment than my own. Once again, this case does not present even a *prima facie* problem for competent dogmatism. The alien BIV's priors were fixed causally by interactions with the environment that its species naturally inhabits, its perceptual system is competent and thus its perceptual beliefs are warranted. If, on the other hand, the encoded probability distribution has shifted based on interactions with the artificial world, then it is this latter environment with regards to which the organism's perceptual system must be evaluated.

To generate a case of an introspectively indistinguishable BIV that competent dogmatism *does not* regard as forming warranted perceptual beliefs, we must stipulate a subject in possession of an incompetent perceptual system. Consider another alien BIV. This organism is doubly unfortunate. It is the subject of a cruel experiment in which its

brain has been removed from its body and placed in a vat to be forced to have experiences as if it were a human philosophy Ph.D. student. Thus, it is radically deceived about its environment. But alas, even if it were removed from the vat, its species is particularly bad at perceiving its native environment. Whatever the evolutionary pressures that shaped the development of its ancestors, the result is a perceptual system that systematically misperceives its normal environment due to either a faulty fixing of its priors, or to a radically non-Bayesian method of updating its priors based on some sensory stimulation. Thus, even if the organism were *released* from the vat and placed back in its body, it would still end up being radically deceived about the world. While in the vat, however this organism enjoys experiences that are introspectively indistinguishable from my own. To construct a case we have stipulated an alien brain with some very peculiar features: (1) it is radically different from a human brain, (2) it evolved in such a way as to be perceptually incompetent in its normal environment, and (3) it is wired in such a way that it is capable of being stimulated to have experiences introspectively indistinguishable to my own. What should we say about the epistemic status of this organism's perceptual beliefs?

If the competent dogmatist theory I have developed here is correct, it indicates that we should have reservations about attributing warranted perceptual beliefs to this subject. Such doubts seem to be perfectly appropriate. Here I can only report on my own intuitions, which I find to be quite nebulous. This organism is so different from me both in the way that its cognitive systems work and in the way that it relates to its natural environment that I am inclined to be agnostic about its epistemic status. Could non-envatted members of this alien species form immediately warranted beliefs about their normal environment? Could

they come to know about features of their local environment on the basis of their perceptual experiences alone? If they could, then the explanation for this ability would have to be very different than what explains *our* perceptual warrant. Perhaps one could spell out the further details of the case in such a way that would make it clear that while they are not perceptually competent their perceptual systems or more general cognitive systems have good-making features that support the attribution of warranted beliefs. However, lacking such a specification of the case, I conclude that there is no *prima facie* problem for competent dogmatism presented by this case.

I should emphasize that I do not mean anything in this section to be taken as a response to external world skepticism. In the present context, I am neutral as to whether or not a competent dogmatist can provide a convincing, non-question-begging response to the skeptic. The preceding discussion was meant solely to illustrate that competent dogmatism does not run seriously afoul of intuitions concerning subjects in skeptical scenarios (as do some other externalist theories of warrant). I think it clear that competent dogmatism entails plausible judgments about radically deceived subjects in the many of these cases. Competent dogmatism is consistent with the view that some subjects with experiences introspectively indistinguishable to my own are not warranted in their perceptual beliefs. In these cases, however, I think our intuitive judgments are less clear. Such subjects are so cognitively different from us, that I, at least, lack any clear intuitions about the epistemic status of their beliefs, and would be reticent to base much of substance on such intuitions even if I did possess them.

9. Conclusion

Allow me to conclude this dissertation by considering some avenues of further investigation for the approach to explaining perceptual warrant that I have championed here. Despite my arguments in Chapter II that dogmatism as a general thesis can avoid many of the objections that have been brought to bear against it, there are many difficult questions that must still be addressed. These are no less pressing for the specific version of the position that I have developed in this final chapter. One major set of issues that must be addressed by the view defended here involves the role that immediately warranted perceptual beliefs play in warranting further beliefs. In Chapter II, I argued that the dogmatist is not explicitly committed to any particular answers to such questions. However, a complete story of the epistemic role of perceptual experiences must (eventually) address a host of issues involving the more general epistemic features of believers. Some such questions are: what are the conditions that must be met in order for immediately warranted perceptual beliefs to ground warrant preserving inferences to further beliefs? To what degree, if any, does dogmatism in general and competent dogmatism in particular aid in addressing the challenges of external world skepticism? Does immediate perceptual warrant constitute a sort of *general* foundation (in a broadly Cartesian sense) for our non-basic beliefs? Though questions of this kind are beyond the bounds of the discussion that I have carried out here, I take them to be of great importance. A more complete dogmatist epistemology than I have presented here ought to address them, and I hope to do so in future work.

Throughout this dissertation I have restricted my investigations to *basic perceptual beliefs*, those that involve a relatively narrow range of perceptual contents, while eschewing

discussion of perceptual beliefs more generally. There are several pressing issues here. One has to do with the *cognitive penetrability* of perceptual experience by the background beliefs of a subject. There are cases in which a subject's beliefs can effect the contents of her perceptual experience. For instance, one's beliefs about the typical color of bananas can affect the color that one perceives when one looks at a gray-scale image of a banana in the right conditions ((as in Hansen et. al (2006)). The degree to which perceptual experiences are cognitively penetrable in this sense is controversial,[131] a more fully worked out version of the view I have defended here must address such phenomena and the extent to which they can affect the immediacy or the degree to which one's basic perceptual beliefs are warranted.

It would also be beneficial to have something to say about perceptual beliefs more generally. As I said in Chapter I, there are many beliefs that are, intuitively, perceptual beliefs that dogmatism does not address in virtue of its restriction to *basic perceptual beliefs*. For instance, my belief that there is a laptop computer on the table in front of me likely requires some warranted background beliefs about the general appearance and functioning of laptop computers in order for it to be warranted. Nonetheless, the belief is intuitively perceptual. That is, I form the belief on the basis of my perceptual experiences and, in most cases, without reflecting on my background beliefs about laptops or engaging in any conscious inference. Such cases of perceptual belief formation are ubiquitous in the daily life of a normal adult human. A complete account of the role of perception in epistemology must say something about the conditions that must be met for such beliefs to be warranted,

[131] See Lyons (2011), Siegel (2012) and Silins (2013) for recent discussions of these issues

and the relationship of their warrant to the immediate warrant a subject enjoys for her basic perceptual beliefs. Any attempt to do so is made problematic by a host of difficult questions. To what degree are these beliefs inferential? To what extent can the perceptual system (or some closely related subsystem) take background beliefs of a subject as inputs in its processing? Are these background beliefs warranted, and if they are what explains their warrant? How is their warrant related to the warrant we enjoy for our *basic* perceptual beliefs? To what extent, if any, can the conceptual repertoire of a visual system be expanded and supplemented by the higher-level conceptual resources of a subject?

I am tentatively of the opinion that competent dogmatism is in a good position to handle these difficulties. Roughly, the strategy would be to take such background information available to the subject as affecting perceptual (or post-perceptual) processing by providing additional evidence from which the posterior probabilities are computed by the visual system. In some such cases, there may be additional non-perceptual cognitive sub-systems whose inferences and functioning are not reflectively accessible to the subject. Many of the norms associated with the inferential processes of perceptual systems could be applied to the functioning of these processes as well. If these background beliefs must be warranted in order for such penetrated experiences to warrant perceptual beliefs, this may affect the immediacy of *particular* perceptual beliefs, without significantly altering the general framework I have defended here. Of course, this is a mere promissory note for a more complete accounting; but I am optimistic about the prospects of a supplemented version of competent dogmatism to systematically deal with issues of this kind.

Thus, though many pressing issues regarding perceptual warrant remain, the thesis that our basic perceptual beliefs are immediately and defeasibly warranted remains well-supported and undefeated by the various objections marshaled against it. Furthermore, I take it that my arguments in the present chapter have established that competent dogmatism is a plausible theory with explanatory power sufficient to warrant further consideration and development.

References

Adelson, E. (1995). Checkershadow Illusion. Retrieved May 31, 2015, from http://web.mit.edu/persci/people/adelson/checkershadow_illusion.html

Alston, W. P. (1980). Level-Confusions in Epistemology. *Midwest Studies in Philosophy*, *5*(1), 135–150.

Altschul, J. (2012). Entitlement, Justification, and the Bootstrapping Problem. *Acta Analytica*, *27*(4), 345–366.

Azzopardi, P., & Cowey, A. (1998). Blindsight and Visual Awareness. *Consciousness and Cognition*, *7*(3), 292–311.

Becker, K. (2012). Basic Knowledge and Easy Understanding. *Acta Analytica*, *27*(2), 145–161.

Becker, K. (2013). Why Reliabilism Does Not Permit Easy Knowledge. *Synthese*, *190*(17), 3751–3775.

BonJour, L. (1985). *The Structure of Empirical Knowledge*. Harvard University Press.

Bouwsma, O. K. (1949). Descartes' evil genius. *The Philosophical Review*, 141–151.

Brainard, D. (2009). Bayesian approaches to color vision. *The Visual Neurosciences*, *4*.

Brainard, D. H., & Freeman, W. T. (1997). Bayesian color constancy. *JOSA A*, *14*(7), 1393–1411.

Brainard, D. H., Longère, P., Delahunt, P. B., Freeman, W. T., Kraft, J. M., & Xiao, B. (2006). Bayesian model of human color constancy. *Journal of Vision*, *6*(11), 10.

Brogaard, B. (2013). Phenomenal Seemings and Sensible Dogmatism. In C. Tucker (Ed.), *Seemings and Justification* (p. 270). Oxford University Press.

Brueckner, A. (2012). Perceptual Anti-Individualism and Skepticism. *International Journal for the Study of Skepticism*, *2*(2), 145–151.

Brueckner, A. (2013). *Essays on Skepticism* (Reprint edition). Oxford; New York: Oxford University Press.

Burge, T. (1989). Individualism and psychology. In *Rerepresentation* (pp. 39–75). Springer.

Burge, T. (1993). Content preservation. *The Philosophical Review*, 457–488.

Burge, T. (2003). Perceptual Entitlement. *Philosophy and Phenomenological Research*, *67*(3), 503–48.

Burge, T. (2007). *Foundations of Mind: Philosophical Essays, Volume 2* (1 edition). Oxford : Oxford ; New York: Oxford University Press.

Burge, T. (2010). *Origins of Objectivity* (1 edition). Oxford ; New York: Oxford University Press.

Burge, T. (2013). *Cognition Through Understanding: Self-Knowledge, Interlocution, Reasoning, Reflection: Philosophical Essays, Volume 3* (1 edition). Oxford: Oxford University Press.

Chalmers, D. J. (2006). Perception and the Fall From Eden. In T. S. Gendler & J. Hawthorne (Eds.), *Perceptual Experience* (pp. 49–125). Oxford University Press.

Cohen, S. (1984). Justification and Truth. *Philosophical Studies*, *46*(3), 279–95.

Cohen, S. (1988). How to Be a Fallibilist. *Philosophical Perspectives*, *2*, 91–123.

Cohen, S. (1998). Contextualist solutions to epistemological problems: Scepticism, Gettier, and the lottery. *Australasian Journal of Philosophy*, *76*(2), 289–306.

Cohen, S. (2002). Basic Knowledge and the Problem of Easy Knowledge. *Philosophy and Phenomenological Research*, *65*(2), 309–329.

Cohen, S. (2005). Why Basic Knowledge is Easy Knowledge. *Philosophy and Phenomenological Research*, *70*(2), 417–430.

Cohen, S. (2010). Bootstrapping, Defeasible Reasoning, and a Priori Justification. *Philosophical Perspectives*, *24*(1), 141–159.

Conee, E., & Feldman, R. (1998). The generality problem for reliabilism. *Philosophical Studies*, *89*(1), 1–29.

Cowey, A. (2004). The 30th Sir Frederick Bartlett Lecture: Fact, Artifact, and Myth About Blindsight. *Quarterly Journal of Experimental Psychology A*, *57*(4), 577–609.

Cullison, A. (2010). What Are Seemings? *Ratio*, *23*(3), 260–274.

Cullison, A. (2013). Seemings and Semantics. *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*.

Cuneo, T. (2009). Themes From G.E. Moore: New Essays in Epistemology and Ethics • by Susana Nuccetelli and Gary Seay. *Analysis*, *69*(1), 167–169.

Davidson, D. (1987). Knowing one's own mind. In *Proceedings and addresses of the American philosophical association* (pp. 441–458).

DePoe, J. M. (2011). Defeating the self-defeat argument for phenomenal conservativism. *Philosophical Studies*, *152*(3), 347–359.

Dretske, F. (1995). *Naturalizing the Mind*. MIT Press.

Dretske, F. (2003). Experience as Representation. *Philosophical Issues*, *13*(1), 67–82.

Feldman, J. (Forthcoming). Bayesian Models of Perception: A Tutorial Introduction. In *Oxford Handbook of Perceptual Organization*.

Feldman, R., & Conee, E. (1985). Evidentialism. *Philosophical Studies*, *48*(1), 15–34.

Fumerton, R. A. (1995). *Metaepistemology and Skepticism*. Rowman & Littlefield.

Geisler, W. S. (2008). Visual perception and the statistical properties of natural scenes. *Annu. Rev. Psychol.*, *59*, 167–192.

Gettier, E. (1963). Is Justified True Belief Knowledge? *Analysis*, *23*(6), 121–123.

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception: Classic Edition* (1 edition). New York; Hove, England: Psychology Press.

Goldman, A. I. (1967). A Causal Theory of Knowing. *Journal of Philosophy*, *64*(12), 357–372.

Goldman, A. I. (1979). What is justified belief? In *Justification and knowledge* (pp. 1–23). Springer.

Goldman, A. I. (1986). *Epistemology and Cognition*. Harvard University Press.

Goldman, A. I. (1988). Strong and weak justification. *Philosophical Perspectives*, 51–69.

Goldman, A. I. (1992). *Liaisons: Philosophy meets the cognitive and social sciences*. MIT Press.

Graham, P. J. (2011). Perceptual Entitlement and Basic Beliefs. *Philosophical Studies*, *153*(3), 467–475.

Gregory, R. L. (1997). *Eye and Brain* (5th edition). Princeton, N.J: Princeton University Press.

Hahn, M., & Ramberg, B. (Eds.). (2003). *Reflections and Replies: Essays on the Philosophy of Tyler Burge*. Cambridge, Mass: A Bradford Book.

Hales, S. D. (1995). Epistemic Closure Principles. *Southern Journal of Philosophy*, *33*(2), 185–202.

Hansen, T., Olkkonen, M., Walter, S., & Gegenfurtner, K. R. (2006). Memory modulates color appearance. *Nature Neuroscience*, *9*(11), 1367–1368.

Hoffman, D. D. (2009). The user-interface theory of perception: Natural selection drives true perception to swift extinction. Citeseer.

Huemer, M. (2001). *Skepticism and the Veil of Perception*. Lanham: Rowman &Amp; Littlefield.

Huemer, M. (2006). Phenomenal Conservatism and the Internalist Intuition. *American Philosophical Quarterly*, *43*(2), 147–158.

Huemer, M. (2007). Compassionate Phenomenal Conservatism. *Philosophy and Phenomenological Research*, *74*(1), 30–55.

Huemer, M. (2011). Phenomenal Conservatism and Self-Defeat: A Reply to DePoe. *Philosophical Studies*, *156*(1), 1–13.

Huemer, M. (2013a). Epistemological Asymmetries Between Belief and Experience. *Philosophical Studies*, *162*(3), 741–748.

Huemer, M. (2013b). Phenomenal Conservatism Über Alles. In C. Tucker (Ed.), *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism* (p. 328). Oxford University Press

IllusionWorks, L.L.C. (1997). Ames Room. Retrieved July 20, 2015, from http://psylux.psych.tu-dresden.de/i1/kaw/diverses%20Material/www.illusionworks.com/html/ames_room.html

Kentridge, R. W., & Heywood, C. A. (1999). The Status of Blindsight: Near-Threshold Vision, Islands of Cortex and the Riddoch Phenomenon. *Journal of Consciousness Studies*, *6*(5), 3–11.

Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as Bayesian inference. *Annu. Rev. Psychol.*, *55*, 271–304.

Kersten, D., & Yuille, A. (2003). Bayesian models of object perception. *Current Opinion in Neurobiology*, *13*(2), 150–158.

Kitaoka, A. (2003). Rotating Snakes. Retrieved June 10, 2013, from http://www.ritsumei.ac.jp/~akitaoka/rotsnakee.html

Knill, D. C., Kersten, D., & Yuille, A. (1996). Introduction: A Bayesian formulation of visual perception. *Perception as Bayesian Inference*, 1–21.

Knill, D. C., & Richards, W. (1996). *Perception as Bayesian inference*. Cambridge University Press.

Kornblith, H. (2009). A Reliabilist Solution to the Problem of Promiscuous Bootstrapping.

Analysis, *69*(2), 263–267.

Lehrer, K. (1990). *Theory of Knowledge* (1 edition). London: Routledge.

Lyons, J. (2009). *Perception and Basic Beliefs: Zombies, Modules, and the Problem of the External World*. Oxford University Press.

Lyons, J. (2011). Circularity, reliability, and the cognitive penetrability of perception. *Philosophical Issues*, *21*(1), 289–311.

Mack, D. (2013). *Hollow Face Illusion* [Photograph]. Retrieved from http://fineartamerica.com/featured/3-hollow-face-illusionartwork-david-mack.html

Mamassian, P., Landy, M. S., & Maloney, L. T. (2002). Bayesian modelling of visual perception. *Probabilistic Models of the Brain: Perception and Neural Function*, 13–36.

Marcel, A. J. (1998). Blindsight and Shape Perception: Deficit of Visual Consciousness or of Visual Function? *Brain*, *121*, 1565–88.

Markie, P. J. (2005a). Easy Knowledge. *Philosophy and Phenomenological Research*, *70*(2), 406–416.

Markie, P. J. (2005b). The Mystery of Direct Perceptual Justification. *Philosophical Studies*, *126*(3), 347–373.

Markie, P. J. (2013). Searching for True Dogmatism. In C. Tucker (Ed.), *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism* (p. 248). Oxford University Press

Mark, J. T., Marion, B. B., & Hoffman, D. D. (2010). Natural selection and veridical perceptions. *Journal of Theoretical Biology*, *266*(4), 504–515.

Marr, D. (1982). *Vision: A computational approach*. Freeman & Co., San Francisco.

Matthen, M. (Ed.). (2015). *The Oxford Handbook of the Philosophy of Perception* (1 edition). Oxford University Press.

McDowell, J. (1996). *Mind and World* (1st Harvard University Press paperback ed edition). Cambridge, Mass: Harvard University Press.

McGrath, M. (2013a). Dogmatism, Underminers and Skepticism. *Philosophy and Phenomenological Research*, *86*(3), 533–562.

McGrath, M. (2013b). Phenomenal Conservatism and Cognitive Penetration: The Bad Basis Counterexamples. In C. Tucker (Ed.), *Seemings and Justification*.

Moore, G. E. (1962a). A Defense of Common Sense. In *Philosophical Papers*. New York, NY: Collier Books.

Moore, G. E. (1962b). Four Forms of Skepticism. In *Philosophical Papers*. New York, NY: Collier Books.

Moore, G. E. (1962c). *Philosophical Papers*. New York, NY: Collier Books.

Moore, G. E. (1962d). Proof of an External World. In *Philosophical Papers*. New York, NY: Collier Books.

Moretti, L. (forthcoming). In Defence of Dogmatism. *Philosophical Studies*, 1–22.

Orlandi, N. (2014). *The Innocent Eye: Why Vision Is Not a Cognitive Process*. Oxford University Press.

Palmer, S. E. (1999). *Vision Science: Photons to Phenomenology* (1 edition). Cambridge, Mass: The MIT Press.

Peacocke, C. (2004a). Explaining Perceptual Entitlement. In R. Schantz (Ed.), *The Externalist Challenge* (pp. 441–80). De Gruyter.

Peacocke, C. (2004b). *The Realm of Reason*. Oxford University Press.

Persaud, N., & Cowey, A. (2008). Blindsight is Unlike Normal Conscious Vision: Evidence From an Exclusion Task. *Consciousness and Cognition*, *17*(3), 1050–1055.

Plantinga, A. (1993). *Warrant and Proper Function*. Oxford University Press.

Pryor, J. (2000). The Skeptic and the Dogmatist. *Noûs*, *34*(4), 517–549.

Pryor, J. (2004). What's Wrong with Moore's Argument? *Philosophical Issues*, *14*(1), 349–378.

Pryor, J. (2012). When Warrant Transmits. In C. Wright & A. Coliva (Eds.), *Mind, Meaning, and Knowledge: Themes From the Philosophy of Crispin Wright*. Oxford University Press.

Pryor, J. (2013). Problems for Credulism. In C. Tucker (Ed.), *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*.

Putnam, H. (1981). Brains in a vat. In *Reason, Truth and History* (pp. 1–21). Cambridge, MA: Cambridge University Press.

Rao, R. P., Olshausen, B. A., & Lewicki, M. S. (2002). *Probabilistic models of the brain: Perception and neural function*. MIT press.

Rescorla, M. (2013). Bayesian perceptual psychology. *The Oxford Handbook of the Philosophy of Perception. Oxford*.

Rock, I. (1985). *Perception* (1st edition). New York: Scientific American Books - W. H. Freeman & Co.

Schiffer, S. (2004). Skepticism and the vagaries of justified belief. *Philosophical Studies*, *119*(1), 161–184.

Siegel, S. (2012). Cognitive penetrability and perceptual justification*. *Nous*, *46*(2), 201–222.

Siegel, S., & Silins, N. (forthcominga). Attention and Perceptual Justification. In A. Pautz & D. Stoljar (Eds.), *Festschrift for Ned Block*. MIT Press.

Siegel, S., & Silins, N. (forthcomingb). The Epistemology of Perception. In M. Matthen (Ed.), *Oxford Handbook of Philosophy of Perception*. Oxford.

Silins, N. (2007). Basic Justification and the Moorean Response to the Skeptic. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford Studies in Epistemology: Volume 2*. Oxford University Press.

Silins, N. (2008). Basic Justification and the Moorean Response to the Skeptic. In *Oxford Studies in Epistemology Volume 2*. Oxford University Press.

Silins, N. (2011). Seeing Through the "Veil of Perception." *Mind*, *120*(478), 329–367.

Silins, N. (2012). Explaining Perceptual Entitlement. *Erkenntnis*, *76*(2), 243–261.

Silins, N. (2013). The significance of high-level content. *Philosophical Studies*, *162*(1), 13–33.

Silvanto, J., Cowey, A., Lavie, N., & Walsh, V. (2007). Making the blindsighted see. *Neuropsychologia*, *45*(14), 3346–3350.

Smithies, D. (forthcoming). Perception and the External World. *Philosophical Studies*.

Smithies, D. (2014). The Phenomenal Basis of Epistemic Justification. In J. Kallestrup & M. Sprevak (Eds.), *New Waves in Philosophy of Mind* (pp. 98–124). Palgrave MacMillan.

Sosa, E. (2009). *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford University Press.

Sosa, E. (2011). *Reflective Knowledge: Apt Belief and Reflective Knowledge, Volume II*. Oxford University Press.

Steup, M. (2013). Does Phenomenal Conservatism Solve Internalism's Dilemma? In C. Tucker (Ed.), *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism* (p. 135). Oxford University Press.

Stoerig, P., & Barth, E. (2001). Low-Level Phenomenal Vision Despite Unilateral Destruction of Primary Visual Cortex. *Consciousness and Cognition*, *10*(4), 574–587.

Stoerig, P., & Cowey, A. (1989). Wavelength Sensitivity in Blindsight. *Nature*, *342*, 916–18.

Storkey, A. Amos Storkey: Hompage [Homepage]. Retrieved May 31, 2015, from http://homepages.inf.ed.ac.uk/amos/craterorhill.html

Stroud, B. (1984). *The significance of philosophical scepticism*. Oxford University Press.

Sun, J., & Perona, P. (1998). Where is the sun? *Nature Neuroscience*, *1*(3), 183–184.

Tolhurst, W. E. (1998). Seemings. *American Philosophical Quarterly*, *35*(3), 293–302.

Tooley, M. (2013). Michael Huemer and the Principle of Phenomenal Conservatism. In C. Tucker (Ed.), *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism* (p. 306). Oxford University Press.

Tucker, C. (2013). *Seemings and justification: New essays on dogmatism and phenomenal conservatism*. Oxford University Press.

Tucker, C. (2013). Seemings and Justification: An Introduction. In C. Tucker (Ed.), *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*. Oxford University Press.

Tucker, C. (2010). Why Open-Minded People Should Endorse Dogmatism. *Philosophical Perspectives*, *24*(1), 529–545.

Tye, M. (1993). Qualia, Content, and the Inverted Spectrum. *Noûs*, *27*(2), 159–183.

Unger, P. (1979). *Ignorance: A Case for Scepticism*. New York: Clarendon Press/Oxford.

Van Cleve, J. (2003). Is Knowledge Easy – Or Impossible? Externalism as the Only Alternative to Skepticism. In S. Luper (Ed.), *The Skeptics: Contemporary Essays*. Ashgate.

Vogel, J. (2000). Reliabilism Leveled. *Journal of Philosophy*, *97*(11), 602–623.

Wagemans, J. (Ed.). (2015). *Oxford Handbook of Perceptual Organization*.

Weisberg, J. (2010). Bootstrapping in General. *Philosophy and Phenomenological Research*, *81*(3), 525–548.

Weisberg, J. (2012). The Bootstrapping Problem. *Philosophy Compass*, *7*(9), 597–610.

Weiskrantz, L. (1996). Blindsight Revisited. *Current Opinion in Neurobiology*, *6*, 215–220.

Weiskrantz, L. (1997). *Consciousness Lost and Found*. Oxford University Press.

Weiskrantz, L. (2009). *Blindsight: A Case Study Spanning 35 Years and New Developments*. Oxford University Press Oxford.

Weiskrantz, L., Barbur, J. L., & Sahraie, A. (1995). Parameters Affecting Conscious Versus Unconscious Visual Discrimination Without V. *Proceedings of the National Academy of Sciences Usa*, *92*, 6122–26.

Weiss, Y., Simoncelli, E. P., & Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nature Neuroscience*, *5*(6), 598–604.

White, R. (2006). Problems for Dogmatism. *Philosophical Studies*, *131*(3), 525–57.

Wittgenstein, L. (2003). *Philosophical Investigations: The German Text, with a Revised English Translation*. Malden, Ma, Blackwell Pub.

Woodward, J. (2005). *Making Things Happen: A Theory of Causal Explanation* (1 edition). Oxford: Oxford University Press.

Wright, C. (2004). Wittgensteinian Certainties. In D. McManus (Ed.), *Wittgenstein and Scepticism* (pp. 22–55). Routledge.

Wright, C. (2007). The Perils of Dogmatism. In Nuccetelli & Seay (Eds.), *Themes from G. E. Moore: New Essays in Epistemology*. Oxford University Press.